

VMware Cloud Foundation Operations and Administration Guide

14 JAN 2020

VMware Cloud Foundation 3.9

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2015-2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About the VMware Cloud Foundation Operations and Administration Guide	9
1 Administering Cloud Foundation Systems	11
VMware Software Components Deployed in a Typical Cloud Foundation System	12
Web Interfaces Used When Administering Your Cloud Foundation System	12
2 Getting Started with SDDC Manager	14
Log In to the SDDC Manager Dashboard	14
Tour of the SDDC Manager User Interface	15
Log out of the SDDC Manager Dashboard	19
3 Configuring Customer Experience Improvement Program	20
4 Certificate Management	22
View Certificate Information	23
Configure a Microsoft Certificate Authority	24
Prepare the Certificate Service Template	25
Add OpenSSL CA support	25
Install Certificates with the Microsoft Certificate Authority or Inbuilt OpenSSL Certificate Authority	26
Install Certificates with External or Third-Party Certificate Authorities	29
Clean Out Old or Unused Certificates	32
5 License Management	34
Add License Keys for the Software in Your Cloud Foundation System	34
Edit License Description	35
Enable vRealize Log Insight Logging for Workload Domains	35
Delete License Key	36
6 Installing ESXi Software on Cloud Foundation Servers	37
Download ESXi Software and VIBs	38
Provide Network Information for Imaging	38
Upload ISOs and VIBs to the VMware Imaging Appliance	39
Image Servers with ESXi and VIBs	41
Post-Imaging Tasks	42
7 Host Management	44
About Network Pools	44

Sizing a Network Pool	45
Create a Network Pool	46
View Network Pool Details	47
Edit a Network Pool	47
Delete a Network Pool	48
Commission Hosts	48
Decommission Hosts	51
Cleaning Up Decommissioned Hosts	52
Clean up a Decommissioned Host Using the SoS Utility	52
Clean up a Decommissioned Host Using the Direct Console User Interface	53
View Host Inventory	54
8 Working with the Management Domain and VI Workload Domains	57
Adding Virtual Machines to the Management Domain	59
About VI Workload Domains	60
Prerequisites for a Workload Domain	61
Additional Prerequisites for an NSX-T Based Workload Domain	64
Start the VI Configuration Wizard	66
Deploy and Configure NSX Edges in Cloud Foundation	72
Sample Values for Deploying NSX-T Edges	73
Create Profiles	74
Create Transport Zones	75
Create NSX-T Segments for System, Uplink, and Overlay Traffic	75
Deploy NSX-T Edge Appliances	76
Join the NSX-T Edge Nodes to the Management Plane	79
Configure Edge Nodes as Transport Nodes	80
Create an NSX-T Edge Cluster	81
Configure Dynamic Routing	82
View Workload Domain Details	87
Delete a VI Workload Domain	87
View Cluster Details	88
Shrink a Workload Domain	89
Remove a Host from a Cluster in a Workload Domain	89
Delete a Cluster from a Workload Domain	90
Expand a Workload Domain	91
Add a Host to a Cluster in a Workload Domain	91
Add a Cluster to a Workload Domain	92
9 Working with VMware Enterprise PKS	95
Deploy an Enterprise PKS Solution	96
Prerequisites for Deploying VMware Enterprise PKS	96

- Start the Deploy Wizard and Specify General Settings 99
- Specify NSX-T Settings 100
- Specify PKS Settings 101
- Provide Certificate Details 102
- Specify Management Availability Zones 102
- Specify Kubernetes Availability Zones 103
- Configure Compute Availability Zones 104
- Review Summary 105
- View PKS Solution Details 106
- Delete a Enterprise PKS Solution 108
- Expand or Shrink an NSX-T Workload Domain associated with an PKS Solution 109

- 10 Working with Horizon Domains 110**
 - Sizing Guidelines 113
 - Prerequisites for a Horizon Domain 115
 - Create a Horizon Domain 118
 - Select VI Workload Domains for the Horizon Domain 119
 - Provide Active Directory Details 119
 - Provide SQL Server Details 120
 - Add Load Balancers 121
 - Add Connection Servers 122
 - Add Composer Servers 124
 - Add Unified Access Gateway Appliances 125
 - Add App Volumes 126
 - Add User Environment Manager 127
 - Review Horizon Domain Configuration 128
 - Resume Horizon Domain Creation 129
 - Exporting and Importing a Horizon Domain Configuration 130
 - Export a Horizon Domain Configuration 130
 - Import a Horizon Domain Configuration 131
 - View Horizon Domain Details 137
 - Expand a Horizon Domain 137
 - Delete Horizon Domain 138

- 11 Deploy a Workload Domain or Cluster at a Remote Location 139**
 - Prerequisites for Cloud Foundation Remote Clusters 140
 - Deploy a Workload Domain at a Remote Location 140
 - Add a Cluster at a Remote Location 141

- 12 vRealize Suite Products and Cloud Foundation 143**
 - Deploy vRealize Suite Lifecycle Manager in Cloud Foundation 146

- Adding vRealize Automation to Cloud Foundation 148
 - Add a vRealize Automation License Key to Cloud Foundation 148
 - Deploy vRealize Automation in Cloud Foundation 149
 - Post-Deployment Tasks for vRealize Automation in Cloud Foundation 156
- Adding vRealize Operations to Cloud Foundation 159
 - Add a vRealize Operations License Key to Cloud Foundation 159
 - Deploy vRealize Operations in Cloud Foundation 160
 - Post-Deployment Tasks for vRealize Operations in Cloud Foundation 163
- Connect vRealize Suite Products to Workload Domains in Cloud Foundation 164
 - Connect vRealize Suite Products to Workload Domains in Cloud Foundation 165
 - Connect Workload Domains to vRealize Suite Products in Cloud Foundation 166
 - Enable vRealize Log Insight in Cloud Foundation 167
- Add a Node to a vRealize Operations Analytics Cluster 168
 - Update Apache Configuration for New Nodes 169

- 13 Downloading an Install Bundle 170**

- 14 Multi-Instance Management 171**
 - About the Multi-Instance Management Dashboard 173
 - Create a Federation 176
 - Invite a Cloud Foundation Instance to Join a Federation 178
 - Join a Federation 179
 - Join a Federation by Clicking an Invitation 179
 - Join a Federation through the Multi-Instance Management Dashboard 180
 - Leave a Federation 181
 - Dismantle a Federation 181

- 15 Stretching Clusters 183**
 - About Availability Zones and Regions 183
 - Prerequisites for Stretching a Cluster 184
 - Stretch a Cluster 185
 - Unstretch a Cluster 188
 - Expand a Stretched Cluster 190
 - Replace a Failed Host in a Stretched Cluster 192

- 16 Composability Management 195**
 - Configure Translation Layer 195
 - Compose a Server 197
 - View Composability Information 198
 - Add Storage 198
 - Remove Storage 198

- Decompose a Server 199

- 17 Monitoring Capabilities in the Cloud Foundation System 200**
 - Viewing Tasks and Task Details 201
 - Using vRealize Log Insight Capabilities in Your Cloud Foundation System 202
 - Get Started Using the vRealize Log Insight Instance 203

- 18 Updating Cloud Foundation DNS and NTP Servers 205**
 - Update DNS Server Configuration 205
 - Update NTP Server Configuration 207

- 19 Supportability and Serviceability (SoS) Utility 210**
 - SoS Utility Options 210
 - Collect Logs for Your Cloud Foundation System 216
 - Component Log Files Collected By the SoS Utility 218

- 20 Managing Shutdown and Startup of Cloud Foundation 221**
 - Shut Down a Cloud Foundation System 221
 - Start Up a Cloud Foundation System 225

- 21 Replace Host Components 229**
 - Replacing Components of a Host Running in Degraded Mode 229
 - Replace Components of a Workload Domain Host Running in Degraded Mode 229
 - Replace Components of an Unassigned Host Running in Degraded Mode 230
 - Replace a Dead Host 231
 - Replace Boot Disk on a Host 231

- 22 User and Group Management 233**
 - Assign Cloud Foundation Role to AD Users or Groups 233
 - View Role Details 234
 - Remove Cloud Foundation Role for a User or Group 234

- 23 Password Management 235**
 - Configure Dual Authentication 235
 - Rotate Passwords for Managed Entities 236
 - Manually Update Passwords 238
 - Look Up Account Credentials 239
 - Password Management cURL API Reference 240
 - Updating SDDC Manager Passwords 242
 - Update SDDC Manager Root and Super User Passwords 242
 - Update SDDC Manager REST API Account Password 243

[Update Expired SDDC Manager root Password](#) 243

24 [Backing Up and Restoring SDDC Manager](#) 245

[Image-Based Backup and Restore](#) 246

[File-Based Backup and Restore](#) 246

[Configure an External SFTP Server for NSX Manager Backups](#) 246

[Backup SDDC Manager](#) 248

[Restore SDDC Manager](#) 248

25 [Cloud Foundation Glossary](#) 249

About the VMware Cloud Foundation Operations and Administration Guide

The *VMware Cloud Foundation Operations and Administration Guide* provides information about managing a VMware Cloud Foundation™ system, including managing the system's virtual infrastructure, managing users, configuring and deploying service offerings, and upgrading and monitoring the system.

Intended Audience

The *VMware Cloud Foundation Operations and Administration Guide* is intended for cloud architects, infrastructure administrators, and cloud administrators who are familiar with and want to use VMware software to quickly deploy and manage an SDDC. The information in this document is written for experienced data center system administrators who are familiar with:

- Concepts of virtualization, software-defined data centers, virtual infrastructure (VI), and virtual desktop infrastructure (VDI)
- VMware virtualization technologies, such as VMware ESXi™, the hypervisor
- Software-defined networking using VMware NSX®
- Software-defined storage using VMware vSAN™
- IP networks

Additionally, you should be familiar with these VMware software products, software components, and their features:

- VMware vSphere®
- VMware vCenter Server® and VMware vCenter Server® Appliance™
- VMware Platform Services Controller™
- VMware vRealize® Log Insight™
- VMware Horizon®
- VMware App Volumes™

Related Publications

The *VMware Cloud Foundation Planning and Preparation Guide* provides detailed information about the software, tools, and external services that are required for Cloud Foundation.

The *VMware Cloud Foundation Architecture and Deployment Guide* contains detailed information about a Cloud Foundation system, its components, and the network topology of a deployed system.

Administering Cloud Foundation Systems

1

As an SDDC administrator, you use the information in the *VMware Cloud Foundation Operations and Administration* document to understand how to administer and operate your installed Cloud Foundation system.

An administrator of a Cloud Foundation system performs tasks such as:

- Manage certificates and passwords.
- Add capacity to your system.
- Configure and provision the systems and the workload domains that are used to provide service offerings.
- Manage provisioned workload domains.
- Monitor alerts and the health of the system.
- Troubleshoot issues and prevent problems across the physical and virtual infrastructure.
- Perform life cycle management on the Cloud Foundation software components.

Note Perform all Cloud Foundation operations in the SDDC Manager UI. Do not use the vSphere Client/Web Client or VMware Host Client to modify or delete resources which Cloud Foundation has deployed and configured, unless specifically instructed to do so in the Cloud Foundation documentation.

See the *Introducing VMware Cloud Foundation* document for a high-level overview of the Cloud Foundation product and the *VMware Cloud Foundation Deployment Guide* for information on deploying the product.

This chapter includes the following topics:

- [VMware Software Components Deployed in a Typical Cloud Foundation System](#)
- [Web Interfaces Used When Administering Your Cloud Foundation System](#)

VMware Software Components Deployed in a Typical Cloud Foundation System

In a typical Cloud Foundation system, you will encounter specific VMware software that SDDC Manager deploys in the system.

Note For information about which specific editions of each VMware product are licensed for use with the Cloud Foundation license, use the information resources at the Cloud Foundation product information page at <http://www.vmware.com/products/cloud-foundation.html>.

For the exact version numbers of the VMware products that you might see in your Cloud Foundation system after the initial bring-up process, see the *Release Notes* document for your Cloud Foundation version. If the system has been updated after the initial bring-up process using the Life Cycle Management features, see "View Upgrade History" in the *VMware Cloud Foundation Upgrade Guide* for details on how to view the versions of the VMware software components that are within your system.

Caution Do not manually change any of the settings that SDDC Manager sets automatically. If you change the generated settings, like names of VMs, unpredictable results might occur. Do not change settings for the resources that are automatically created and deployed during workflows, the workload domain processes, assigned IP addresses or names, and so on.

You can find the documentation for the following VMware software products and components at docs.vmware.com:

- vSphere (vCenter Server, Platform Services Controller, and ESXi)
- vSAN
- NSX for vSphere
- vRealize Log Insight
- vRealize Operations
- vRealize Automation

Web Interfaces Used When Administering Your Cloud Foundation System

You use SDDC Manager loaded in a browser for the single-point-of-control management of your Cloud Foundation system. This user interface provides centralized access to and an integrated view of the physical and virtual infrastructure of your system.

In addition to using the SDDC Manager Dashboard, you can use the following user interfaces for administration tasks involving their associated VMware software components that are part of a VMware SDDC. All these interfaces run in a browser, and you can launch them from within the SDDC Manager Dashboard.

Launch links are typically identified in the user interface by the launch icon: .

VMware SDDC Web Interface	Description	Launch Link Location in SDDC Manager Dashboard
vSphere	This interface provides direct management of resources managed by the vCenter Server instances, for identity management, and for management of the NSX resources that provide the software-defined networking capabilities of the SDDC. You can also manage object level storage policies for distributed software-defined storage provided by vSAN.	<ol style="list-style-type: none"> 1 On the SDDC Manager Dashboard, click Inventory > Workload Domains. 2 Click View Details for a workload domain. 3 In the Domain column, click the domain name. 4 Click the Services tab.
vRealize Log Insight	When the vRealize Log Insight instance is licensed for use in the system, this interface provides direct access to the logs and event data collected and aggregated in vRealize Log Insight for troubleshooting, trend analysis, and reporting.	<ol style="list-style-type: none"> 5 Click the appropriate launch link.
Platform Services Controllers	Launches the web interface for the selected Platform Services Controller.	
NSX Manager	Launches the NSX Manager web interface.	

Getting Started with SDDC Manager

2

You use SDDC Manager to perform administration tasks on your Cloud Foundation system. This user interface provides an integrated view of the physical and virtual infrastructure and centralized access to manage the physical and logical resources.

You work with the SDDC Manager Dashboard by loading it in a web browser. For the list of supported browsers and versions, see the *Release Notes*.

Note When performing out-of-band (OOB) troubleshooting of hardware, some vendors may use Java-based consoles. Refer to the vendor documentation for supported browsers.

This chapter includes the following topics:

- [Log In to the SDDC Manager Dashboard](#)
- [Tour of the SDDC Manager User Interface](#)
- [Log out of the SDDC Manager Dashboard](#)

Log In to the SDDC Manager Dashboard

You access SDDC Manager through the SDDC Manager Dashboard in a supported browser.

Prerequisites

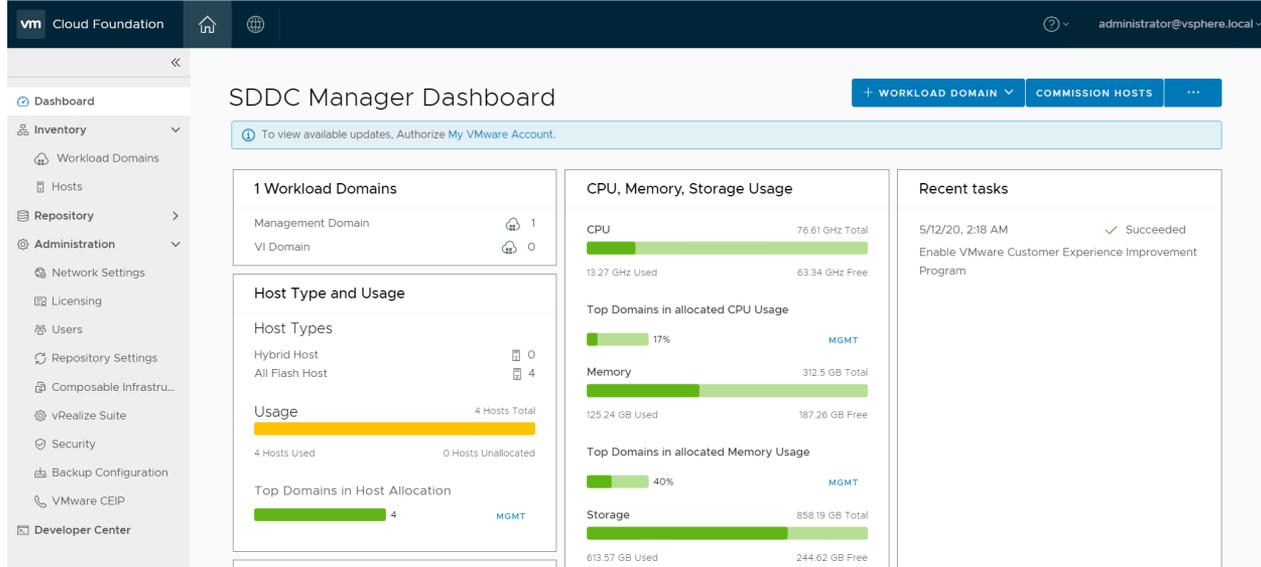
To log in, you need the SDDC Manager IP address or FQDN and the password for the single-sign on user (for example **administrator@vsphere.local**). You added this information to the deployment parameter workbook before bring-up.

Procedure

- 1 In a browser, type one of the following:
 - `https://FQDN` where *FQDN* is the host name of the SDDC Manager.
 - `https://IP_address` where *IP_address* is the IP address of the SDDC Manager.
- 2 Log in with the single-sign on user credentials.

Results

You are logged in to the SDDC Manager and the SDDC Manager Dashboard page appears in the browser.



Tour of the SDDC Manager User Interface

SDDC Manager provides the user interface for your single point of control for managing and monitoring your Cloud Foundation system and for provisioning virtual environments.

You use the Navigation bar to move between the main areas of the user interface.

Navigation Bar

On the left side of the interface is the Navigation bar. The Navigation bar provides a hierarchy for navigating to the corresponding pages.

Category	Functional Areas
Dashboard	<p>The Dashboard provides the high-level administrative view for SDDC Manager features and functions in the form of widgets, including: Workload Domains; CPU, Memory, Storage Usage; Host Types and Usage; Recent Tasks; Ongoing and Scheduled Updates; and Update History.</p> <p>You can control which widgets display and how they are arranged on dashboard.</p> <ul style="list-style-type: none"> ■ To rearrange the widgets, click the heading of the widget and drag it into the desired position. ■ To hide a widget, hover the mouse anywhere over the widget to reveal the X in the upper-right corner, and click the X. ■ To add a widget to the dashboard, click the three dots adjacent to the Commission Hosts button in the upper right corner of the page and select Add New Widgets. This displays all hidden widgets and enables you to select them.
Inventory	<p>The Inventory category directs you to the following destinations:</p> <ul style="list-style-type: none"> ■ Click Workload Domains to go directly to the Workload Domains page, which displays and provides access to all current workload domains and controls for managing workload domains. <p>This page includes detailed status and information about all existing workload domains, including IP addresses, health status, owner, number of hosts, update status, and more. It also displays CPU, memory, and storage utilization for each workload domain, and collectively across all domains.</p> <ul style="list-style-type: none"> ■ Click Hosts to go directly to the Hosts page, which displays and provides access to all current hosts and controls for managing hosts. <p>This page includes detailed status and information about all existing hosts, including IP addresses, network pool, health status, domain and cluster assignment, and storage type. It also displays CPU, memory, and storage utilization for each host, and collectively across all hosts.</p>

Category	Functional Areas
Repository	<p>The Repository category directs you to the following destinations:</p> <ul style="list-style-type: none">■ Click Bundles to view the Cloud Foundation product bundles that are available in the Cloud Foundation instance.■ Click Download History to view the history of update bundle downloads, including version number, date, and other release details. If a bundle is available but has not yet been downloaded, controls for immediate or scheduled downloading appear next to the bundle. <hr/> <p>Note To access patches and bundles, you must be logged in to your myvware account through the Administration > Update Management page.</p>

Category	Functional Areas
Administration	<p>The Administration category directs you to the following destinations:</p> <ul style="list-style-type: none"> ■ Click Network Settings to view and manage network pool settings, including network pool configuration. You can create new pools, and view and modify existing pools. A network pool is a collection of network information with an IP inclusion range reserved for Cloud Foundation. See About Network Pools for more information. ■ Click Licensing to manage VMware product licenses. Add the licenses for the component products that comprise your Cloud Foundation deployment. See Chapter 5 License Management for more information. ■ Click Users to manage Cloud Foundation users and groups, including creating users and groups, setting privileges, assigning roles, and deleting users and groups. See Chapter 22 User and Group Management for more information. ■ Click Repository Settings to log in to your My VMware account, and gain access to install, patch and update bundles. ■ Click Composable Infrastructure to configure composable servers to meet the needs of your workloads without physically moving any hardware components. ■ Click vRealize Suite to deploy and manage vRealize Suite Lifecycle Manager, vRealize Log Insight, vRealize Operations, and vRealize Automation as components of Cloud Foundation. See Chapter 12 vRealize Suite Products and Cloud Foundation for more information. ■ Click Security to configure your certificate authorities and manage password for the accounts that are used by your Cloud Foundation system. See Chapter 4 Certificate Management and Chapter 23 Password Management for more information. ■ Click Backup Configuration to register an external SFTP server with SDDC Manager for backing up NSX Managers. ■ Click VMware CEIP to enroll in the VMware Customer Experience Improvement Program. This program provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, VMware collects technical information about your organization's use of VMware products and services on a regular basis in association with your organization's VMware license key(s).
Developer Center	<p>The Developer Center includes the following sections:</p>

Category	Functional Areas
	<ul style="list-style-type: none"><li data-bbox="810 226 1412 317">■ Overview: API reference documentation. Includes information and steps for all the Public APIs supported by Cloud Foundation.<li data-bbox="810 327 1394 386">■ API Explorer: Lists the APIs and allows you to invoke them directly on your Cloud Foundation system.<li data-bbox="810 396 1350 455">■ Code Samples: Sample code to manage a Cloud Foundation instance.

Log out of the SDDC Manager Dashboard

Log out of SDDC Manager when you have completed your tasks.

Procedure

- 1 In the SDDC Manager Dashboard, open the logged-in account menu by clicking the down arrow next to the account name in the upper right corner.
- 2 Click the menu choice to log out.

Configuring Customer Experience Improvement Program

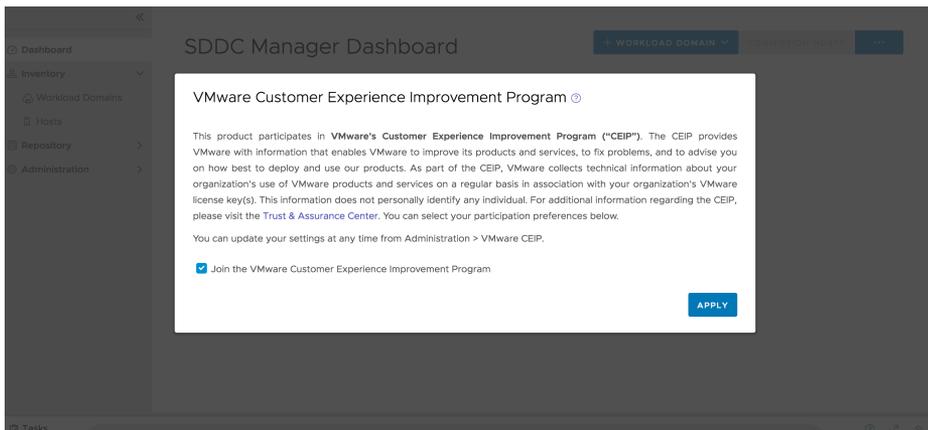
3

This product participates in VMware Customer Experience Improvement Program (CEIP).

The CEIP provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, VMware collects the technical information about your organization's use of the VMware products and services regularly in association with your organization's VMware license keys. This information does not personally identify any individual. For additional information regarding the CEIP, refer to the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

You can enable or disable CEIP across all the SDDC components deployed in VMware Cloud Foundation by the following methods:

- When you log into SDDC Manager for the first time, a pop-up window appears. The **Join the VMware Customer Experience Program** option is selected by default. Deselect this option if you do not want to enable CEIP. Click **Apply**.



- You can enable or disable CEIP from the Administration tab on the SDDC Manager dashboard.

Note When you join CEIP, the **Enable VMware Customer Experience Improvement Program** task is shown in the floating task bar of the SDDC dashboard. Similarly, when you leave CEIP, the **Disable VMware Customer Experience Improvement Program** task is displayed. This task bar is available on all the screens of VMware Cloud Foundation.

To enable or disable CEIP from the **Administration** tab, perform the following steps:

Procedure

- 1** On the SDDC Manager Dashboard, click **Administration > VMware CEIP**.
- 2** To enable CEIP, select the **Join the VMware Customer Experience Improve Program** option.
- 3** To disable CEIP, deselect the **Join the VMware Customer Experience Improve Program** option.

Certificate Management

4

You can manage certificates for all external-facing Cloud Foundation component resources, including configuring a certificate authority, generating and downloading CSRs, and installing them. This section provides instructions for using both Microsoft and non-Microsoft certificate authorities.

You can manage the certificates for the following components.

- Platform Services Controllers
- vCenter Server
- NSX Manager
- SDDC Manager
- vRealize Automation
- vRealize Log Insight
- vRealize Operations

You replace certificates for the following reasons:

- Certificate has expired or is close to expiring.
- Certificate has been revoked.
- You do not want to use the default VMCA certificate.
- Optionally, when you create a new workload domain.

However, it is recommended that you replace all certificates right after deploying Cloud Foundation. After you create new workload domains, you can replace certificates for the appropriate components as needed.

Procedure

1 [View Certificate Information](#)

You can view details of a currently active certificate for a component resource directly in the SDDC Manager Dashboard.

2 Configure a Microsoft Certificate Authority

Before you can generate and install certificates, you must configure a certificate authority (CA).

3 Install Certificates with the Microsoft Certificate Authority or Inbuilt OpenSSL Certificate Authority

You can generate a CSR and signed certificates, and install them for selected resource components directly in the SDDC Manager Dashboard.

4 Install Certificates with External or Third-Party Certificate Authorities

If you intend to generate and install external or third-party certificates, you must download the certificate signing request (CSR) from the SDDC Manager Dashboard and have it manually signed by a third-party CA. You can then use the controls in the SDDC Manager Dashboard to install the certificate.

5 Clean Out Old or Unused Certificates

Old or unused certificates are stored in a trust store in SDDC Manager. You can delete old certificates through the SDDC Manager VM.

View Certificate Information

You can view details of a currently active certificate for a component resource directly in the SDDC Manager Dashboard.

Procedure

1 In the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

2 In the list of domains, click the name of the workload domain to open the details page for that domain.

The workload domain details page displays CPU, memory, and storage allocated to the domain.

3 Select the **Security Tab**.

This tab lists the certificates for each Cloud Foundation resource component, including the following details:

- Issuer, such as Certificate Authority.
- Start and finish dates for certificate validity.
- Current certificate status: Active, Expiring (will expire within 15 days), or Expired.
- Certificate operation status.

- 4 To view certificate details, expand the resource to view the certificate details in the Resource Type column.

The expanded field displays certificate details including signature algorithm, public key, public key algorithm, certificate string, and more.

Configure a Microsoft Certificate Authority

Before you can generate and install certificates, you must configure a certificate authority (CA).

Prerequisites

- Verify that the Microsoft Certificate Authority Server has the correct roles installed. See [Install Microsoft Certificate Authority Roles](#).
- Verify the Microsoft Certificate Authority Server has been configured for basic authentication. See [Configure the Microsoft Certificate Authority for Basic Authentication](#).
- Verify a valid certificate template has been configured on the Microsoft Certificate Authority. See [Create and Add a Microsoft Certificate Authority Template](#).
- Verify least privileged service account has been configured on the Microsoft Certificate Authority Server and Template. See [Assign Certificate Management Privileges to the SDDC Manager Service Account](#).

Note If the CA Web server and CA are on different machines, you must perform the steps mentioned in <https://blogs.technet.microsoft.com/askds/2009/04/22/how-to-configure-the-windows-server-2008-ca-web-enrollment-proxy/> in addition to the following steps.

Procedure

- 1 Navigate to **Administration > Security > Certificate Management** to open the Configure Certificate Authority page.
- 2 Click **Edit** and complete the following configuration settings.

Option	Description
Certificate Authority	Select the CA from the drop-down menu. The default is Microsoft .
CA Server URL	Specify the URL for the CA address server. This address must begin with https:// and end with certsrv , for example https://www.mymicrosoftca.com/certsrv
Username	Provide a valid user name to enable access to the address server.
Password	Provide a valid password to enable access to the address server.
Template Name	Enter the certsrv template name. You must create this template in Microsoft Certificate Authority.

- 3 Click **Save**.

A dialog box appears, asking you to review and confirm the CA server certificate details.

- 4 Click **Accept** to complete the configuration.

Results

The Microsoft CA is now available for use in generating and installing a certificate.

Prepare the Certificate Service Template

To ensure that Cloud Foundation can successfully pass authentication when replacing certificates, you must create the certificate service template with the proper basic authentication configuration through the IIS manager.

Procedure

- 1 Create a Microsoft Active Directory CA with the following features and settings.
 - a Navigate to **Select server roles**.
 - b Under **Active Director Certificate Services**, select **Certification Authority** and **Certification Authority Web Enrollment**.
 - c Under **Web Server (IIS) > Web Server > Security**, select **Basic Authentication**.
- 2 Configure and issue a VMware Certificate Template for **Machine SSL and Solution User certificates** on this CA server.

For step by step procedures, see Knowledge Base article 2112009 [Search Creating a Microsoft Certificate Authority Template for SSL certificate creation in vSphere 6.x](#) .

- 3 Configure the certificate service template and all sites (including default web site) for basic authentication.
 - a Access the IIS manager and navigate to **Server > Sites > Default Web Site > CertSrv**.
 - b Select the Authentication property in the IIS header.
 - c Select and enable **Basic Authentication**.
 - d Restart the site.

What to do next

Use this template when configuring the certificate authority in [Configure a Microsoft Certificate Authority](#).

Add OpenSSL CA support

To generate OpenSSL Certificate Authority (CA) signed certificates for the VMware Cloud Foundation environment:

Procedure

- 1 To configure the OpenSSL CA settings before generating the certificates, navigate to **Administration > Security > Certificate Management**.

- In the Configure Certificate Authority page, select **OpenSSL** for **Certificate Authority**. Provide the required information.

Attribute	Description
Common Name	Specify the FQDN of OpenSSL CA.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization	Specify the name under which your company is legally registered. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Specify the city or the locality where your company is legally registered.
State	Enter the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Select the country where your company is legally registered.

Click **Save**.

- To generate the OpenSSL CA signed certificates, navigate to **Inventory > Workload Domains > Select Domain**.
- Under the **Security** tab, click **Generate Signed Certificates**.
- The **Generate Signed Certificates** pop-up appears. Select **OpenSSL** as the Certificate Authority.
- Click **Generate Certificates**.

Install Certificates with the Microsoft Certificate Authority or Inbuilt OpenSSL Certificate Authority

You can generate a CSR and signed certificates, and install them for selected resource components directly in the SDDC Manager Dashboard.

Prerequisites

- Verify that the bring-up process is complete and successful.
- Verify that you have configured the Certificate Authority, as described in [Configure a Microsoft Certificate Authority](#).

Procedure

- In the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
The Workload Domains page displays information for all workload domains.

- 2 In the list of domains, click the name of the workload domain to open the details page for that domain.

The workload domain details page displays CPU, memory, and storage allocated to the domain.

- 3 Select the **Security Tab**.

This tab lists the default certificates, among other details, for the Cloud Foundation resource components. It also provides controls for working with certificates.

Note You can view the current certificate and key information for a component by clicking the down-arrow icon next to the name.

- 4 Generate the CSR.

- a Use the check boxes to select the resource components for which you want to generate the CSR.

- b Click **Generate CSR**.

The Generate CSRs dialog box opens.

- c Configure the following settings for the CSR.

Option	Description
Algorithm	Select the key type for the certificate. RSA (the default) is typically used. The key type defines the encryption algorithm for communication between the hosts.
Key Size	Select the key size (2048, 3072 or 4096 bit) from the dropdown list.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization	Type name under which your company is legally registered. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State or Province Name	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d Click **Generate CSR**.

The Generate CSRs dialog box closes. The Security tab displays a status of CSR **Generation** is **in progress**. When the CSR generation completes, the **Generate Signed Certificates** button becomes active.

- 5 Generate the signed certificates.
 - a Leave all the resource components selected.
 - b Click **Generate Signed Certificates**.

The Generate Signed Certificates dialog box appears, listing the selected components.

- c For the Select Certificate Authority, select the desired authority, and click **Generate Certificate**.

The Generate Signed Certificates dialog box closes. The Security tab displays a status of **Certificates Generation is in progress**. When the certificate generation completes, the **Install Certificates** button becomes active.

- 6 Click **Install Certificates**.

The Security tab displays a status of **Certificates Installation is in progress**.

Note As installation completes, the Certificates Installation Status column for each selected resource component in the list changes to **Successful** with a green check mark.

Important If you selected SDDC Manager as one of the resource components, you must manually restart SDDC Manager services to reflect the new certificate and to establish a successful connection between Cloud Foundation services and other resources in the management domain.

Important If you selected vRealize Automation as one of the resource components, you must ensure that the vRealize Automation resource root certificate is trusted by all the vRealize Automation VMs in your deployment.

- 7 Restart all services using the provided `sddcmanager_restart_services.sh` script.

To restart the service:

- a Using SSH, log in to the SDDC Manager VM with the following credentials:
 - Username: **vcf**
 - Password: use the password specified in the deployment parameter sheet
- b Enter **su** to switch to the root user.
- c Run the following command:

```
sh /opt/vmware/vcf/operationsmanager/scripts/cli/sddcmanager_restart_services.sh
```

What to do next

If you have replaced the certificate for the vRealize Operations Manager resource component, you must reconfigure the load balancer node. See [Configure SSL Passthrough for vRealize Operations Manager](#) .

Install Certificates with External or Third-Party Certificate Authorities

If you intend to generate and install external or third-party certificates, you must download the certificate signing request (CSR) from the SDDC Manager Dashboard and have it manually signed by a third-party CA. You can then use the controls in the SDDC Manager Dashboard to install the certificate.

Prerequisites

Verify that you have configured and packaged your certificate authority configuration files in the form of a `.tar.gz` file. The contents of this archive must adhere to the following structure:

- The name of the top-level directory must exactly match the name of the domain as it appears in the list on the **Inventory > Workload Domains** page. For example, `MGMT`.
- The PEM-encoded root CA certificate chain file (`rootca.crt`) must reside inside this top-level directory.

The `rootca.crt` file contains a root certificate authority and can have *N* number of intermediate certificates. The file structure of the `rootca.crt` file must look like the following example:

```
-----BEGIN CERTIFICATE-----
<content>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<content>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<content>
-----END CERTIFICATE-----
```

In the above example, there are two intermediate certificates, `intermediate1` and `intermediate2`, and a root certificate. `intermediate1` must use the certificate issued by `intermediate2` and `intermediate2` must use the certificate issued by Root CA.

- This directory must contain one sub-directory for each component resource.

The name of each sub-directory must exactly match the resource hostname of a corresponding component as it appears in the Resource Hostname column in the **Workload Domains > Security** tab.

For example, `nsxManager.vrack.vsphere.local` and `vcenter-1.vrack.vsphere.local`.
- Each sub-directory must contain a corresponding `.csr` file, whose name must exactly match the resource as it appears in the Resource Type column in the **Workload Domains > Security** tab.

For example, the `nsxManager.vrack.vsphere.local` sub-directory contains the `nsxManager.vrack.vsphere.local.csr` file.

- Each sub-directory must contain a corresponding `.crt` file, whose name must exactly match the resource as it appears in the Resource Type column in the **Workload Domains > Security** tab.

For example, the `nsxManager.vrack.vsphere.local` sub-directory must contain the `nsxManager.vrack.vsphere.local.crt` file.

Note All resource and hostname values can be found in the list on the **Inventory > Workload Domains > Security** tab.

Procedure

- 1 In the SDDC Manager Dashboard, navigate to **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 In the list of domains, click the name of the workload domain to open the details page for that domain.

The workload domain details page displays CPU, memory, and storage allocated to the domain.

- 3 Select the **Security Tab**.

This tab lists the default certificates, among other details, for the Cloud Foundation resource components. It also provides controls for working with certificates.

Note You can view the current certificate and key information for a component by clicking the down-arrow icon next to the name.

- 4 Generate the CSR.

- a Use the check boxes to select the resource components for which you want to generate the CSR.

- b Click **Generate CSR**.

The Generate CSRs dialog box opens.

- c Configure the following settings for the CSR.

Option	Description
Algorithm	Select the key type for the certificate. RSA (the default) is typically used. The key type defines the encryption algorithm for the communication between the hosts.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the dropdown list.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization	Enter name under which your company is legally registered. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Enter the city or locality where your company is legally registered.
State or Province Name	Enter the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Enter the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d Click **Generate CSR**.

The Generate CSRs dialog box closes. The Security tab displays a status of CSR Generation is in progress. When CSR generation is complete, the **Download CSR** button becomes active.

- 5 Click **Download CSR** to download and save the CSR files to the directory structure described in the Prerequisites section above.
- 6 External to the SDDC Manager Dashboard, complete the following tasks:
 - a Verify that the different .csr files have successfully generated and are allocated in the required file structure.
 - b Get the certificate requests signed.
This creates the corresponding .crt files.
 - c Verify that the newly acquired .crt files are correctly named and allocated in the required file structure.
 - d Package the file structure as <domain name>.tar.gz.
- 7 Click **Upload and Install**.
- 8 In the Upload and Install Certificates dialog box, click **Browse** to locate and select the newly created <domain name>.tar.gz file.

After you select the file, the **Upload** button becomes active.

9 Click **Upload**.

When the upload is completed, the **Install Certificate** button becomes active.

10 Click **Install Certificate**.

The Security tab displays a status of Certificates Installation is in progress.

Note As the installation is completed, the Certificates Installation Status column for the affected components in the list changes to Successful with a green check mark.

Important If you selected SDDC Manager as one of the resource components, you must manually restart SDDC Manager services to reflect the new certificate and to establish a successful connection between Cloud Foundation services and other resources in the management domain.

Important If you selected vRealize Automation as one of the resource components, you must ensure that all the vRealize Automation VMs in your deployment trust the vRealize Automation resource root certificate.

11 Restart all services using the provided `sddcmanager_restart_services.sh` script.

To restart the service:

- a Using SSH, log in to the SDDC Manager VM with the following credentials:

Username: **vcf**

Password: use the password provided in the deployment parameter sheet

- b Enter **su** to switch to the root user.
- c Run the following command:

```
sh /opt/vmware/vcf/operationsmanager/scripts/cli/sddcmanager_restart_services.sh
```

What to do next

If you have replaced the certificate for the vRealize Operations Manager resource component, you must reconfigure the load balancer node. See [Configure SSL Passthrough for vRealize Operations Manager](#) .

Clean Out Old or Unused Certificates

Old or unused certificates are stored in a trust store in SDDC Manager. You can delete old certificates through the SDDC Manager VM.

Procedure

- 1 Using SSH, log in to the SDDC Manager VM with the following credentials:

Username: **vcf**

Password: use the password specified in the deployment parameter sheet

- 2 Enter **su** to switch to the root user.
- 3 Change to the `/opt/vmware/vcf/operationsmanager/scripts/cli` directory.

```
cd /opt/vmware/vcf/operationsmanager/scripts/cli
```

- 4 From the `/opt/vmware/vcf/operationsmanager/scripts/cli` directory, use the following script and command to discover the names of the certificates in the trust store.

```
sddcmanager-ssl-util.sh -list
```

- 5 Using the name of the certificate, delete the old or unused certificate.

```
sddcmanager-ssl-util.sh -delete <certificate alias name from list>
```

- 6 (Optional) Clean out root certificates in VMware Endpoint Certificate Store from the Platform Services Controller node.

See [Explore Certificate Stores from the vSphere Client](#) in the vSphere product documentation.

License Management

5

In the deployment parameter sheet you completed before bring-up, you entered license keys for the following components:

- VMware vSphere
- VMware vSAN
- VMware NSX for vSphere
- vCenter
- VMware vRealize Log Insight for the management domain

After bring-up, these license keys appear in the Licensing screen of the SDDC Manager Dashboard.

You must have adequate license units available before you create a VI workload domain, add a host to a cluster, or add a cluster to a workload domain. Add license keys as appropriate before you begin any of these tasks.

This chapter includes the following topics:

- [Add License Keys for the Software in Your Cloud Foundation System](#)
- [Edit License Description](#)
- [Enable vRealize Log Insight Logging for Workload Domains](#)
- [Delete License Key](#)

Add License Keys for the Software in Your Cloud Foundation System

You can add licenses to the Cloud Foundation license inventory.

Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Click **+ License Key**.

- 3 Select the product key for which you are entering a license key.
- 4 Type the license key.
- 5 Type a description for the license.

If you have multiple license keys for a product, the description can help in identifying the license. For example, you may want to use one license for high performance workload domains and the other license for regular workload domains.

- 6 Click **Add**.

Edit License Description

If you have multiple license keys for a product, the description can help in identifying the license. For example, you may want to use one license for high performance workload domains and the other license for regular workload domains.

Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Hover your mouse in the license row that you want to edit.
A set of three dots appear on the left of the product name.
- 3 Click the dots and then click **Edit Description**.
- 4 On the Edit License Key Description window, edit the description as appropriate.
- 5 Click **Save**.

Enable vRealize Log Insight Logging for Workload Domains

During the bring-up process, vRealize Log Insight is deployed and configured to collect logs from the management domain components (vSphere, NSX Manager, and SDDC Manager). To enable logging on VI workload domains, you must provide your own license for vRealize Log Insight. After you enter the license key on the vRealize Log Insight UI and enable logging in Cloud Foundation, workload domains are automatically connected to vRealize Log Insight.

Once logging is enabled for workload domains, you cannot disable this setting.

Procedure

- 1 On the SDDC Manager Dashboard, click navigate to **Administration > vRealize Suite**.
- 2 Click **vRealize Log Insight**.
- 3 Click the **vRealize Log Insight** link.
- 4 Login to vRealize Log Insight with the admin credentials you provided in the deployment parameters sheet before bring-up.
- 5 Navigate to **Administration > Management > License**.

- 6 Click **Add New License**.
- 7 Enter the license key and click **Add License**.
- 8 Verify that the license you added is displayed in the license table and the status is active.
Cloud Foundation connects vRealize Log Insight to workload domains.
- 9 On the SDDC Manager Dashboard, click **Enable** in the Enable Logging for all Workload Domains window.

Results

Cloud Foundation connects the vSphere and NSX components for all existing workload domains to vRealize Log Insight. Workload domains created after enabling logging are automatically connected to vRealize Log Insight.

Delete License Key

Deleting a license key removes the license from the Cloud Foundation license inventory. If the license has been applied to any workload domain, host, or cluster, the license continues to work for them.

Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Hover your mouse in the license row that you want to edit.
A set of three dots appear on the left of the product name.
- 3 Click the dots and then click **Remove Key**.
- 4 On the Remove Key dialog box, click **Remove**.

Results

The license is removed from the Cloud Foundation license inventory

Installing ESXi Software on Cloud Foundation Servers

6

You can use the VMware Imaging Appliance (VIA) included with the Cloud Builder VM to image servers for use in the management domain and VI workload domains.

You can use VIA to image servers prior to adding them to Cloud Foundation as part of the host commissioning process. For information about imaging servers prior to bring-up, see the *VMware Cloud Foundation Architecture and Deployment Guide*.

You must have access to the Cloud Builder VM to use the VMware Imaging Appliance. If you deleted VIA after bring-up, you can redeploy it as described in "Deploy Cloud Foundation Builder VM" in the *VMware Cloud Foundation Architecture and Deployment Guide*.

Server Prerequisites

The servers that you image must meet certain prerequisites:

- PXE Boot is configured as primary boot option
- Install device is configured as the second boot option
- Legacy boot mode configured in BIOS (UEFI boot mode is not supported)
- Servers are in the same L2 domain as the Cloud Builder VM
- Servers are reachable over an untagged VLAN/Network (VLAN ID 0)
- The Cloud Builder VM is deployed on an untagged VLAN/Network
- Server hardware/firmware should be configured for virtualization and vSAN and match the Cloud Foundation BOM as described in the Release Notes
- Physical hardware health status should be "healthy" without any errors
- Any onboard NICs are disabled on the servers and only the two 10 GbE NICs reserved for use with Cloud Foundation are enabled in BIOS

The default root credentials for servers imaged with VIA are user **root**, password **EvoSddc!2016**.

This chapter includes the following topics:

- [Download ESXi Software and VIBs](#)
- [Provide Network Information for Imaging](#)
- [Upload ISOs and VIBs to the VMware Imaging Appliance](#)

- [Image Servers with ESXi and VIBs](#)
- [Post-Imaging Tasks](#)

Download ESXi Software and VIBs

In order to image your servers, you need to download an ESXi ISO and any vSphere Installation Bundles (VIBs) required to get the servers to a supported version of ESXi. See the BOM section of the VMware Cloud Foundation Release Notes for information about ESXi support.

You can download the ISO and VIBs from My VMware (<https://my.vmware.com>) to any location on the Windows machine that is connected to the Cloud Builder VM. Make sure to record the MD5 or SHA-1 checksums. You will need them when you upload the ISO/VIB to the VMware Imaging Appliance.

Provide Network Information for Imaging

You must provide the VMware Imaging Appliance with certain network information specific to your environment before you can image your servers. This information is contained in the `via.properties` file on the Cloud Builder VM.

Procedure

- 1 SSH into the Cloud Builder VM using the credentials specified when you deployed the VM. See "Deploy Cloud Foundation Builder VM" in the *VMware Cloud Foundation Architecture and Deployment Guide*.
- 2 Type `su` to switch to the root user.
- 3 Navigate to the `/opt/vmware/evorack-imaging/config/` directory.

4 Update the `via.properties` file with your network information.

- a If the Cloud Builder VM is using the `eth0` interface (default), then you do not need to modify any of the properties in Section A. If the Cloud Builder VM has multiple network interfaces and is not using `eth0`, you must update the following properties.

Property	Description
<code>via.network.interface</code>	Interface of the Cloud Builder VM configured in management network.
<code>via.web.url</code>	The IP address used to access the VMware Imaging Appliance UI. Update this with the IP address of Cloud Builder VM in the management network.
<code>via.network.ifaceaddr</code>	Update this with the IP address of Cloud Builder VM in the management network.
<code>via.dhcp.esxi.tftpServer</code>	IP address of the server where TFTP is running. Update this with the IP address of Cloud Builder VM in the management network.
<code>via.config.remote.pxe=false</code>	Do not modify.

- b Update Section B with the network information for your environment.

Property	Description
<code>via.dhcp.netmask</code>	Netmask of the management network.
<code>via.dhcp.subnet</code>	Subnet of the management network.
<code>via.dhcp.routers</code>	Gateway IP of the management network.
<code>via.esxi.firewall.allowed.network</code>	CIDR notation for subnet IP of the management network.

5 Type `systemctl restart imaging.service` to restart the imaging service.

Wait for the imaging service to restart.

6 Type `systemctl status imaging.service` to verify that the imaging service is running.

What to do next

Log in to the VMware Imaging Appliance and upload software.

Upload ISOs and VIBs to the VMware Imaging Appliance

After you have downloaded the required software and updated `via.properties` with your network information, you can upload ISOs and VIBs to the VMware Imaging Appliance.

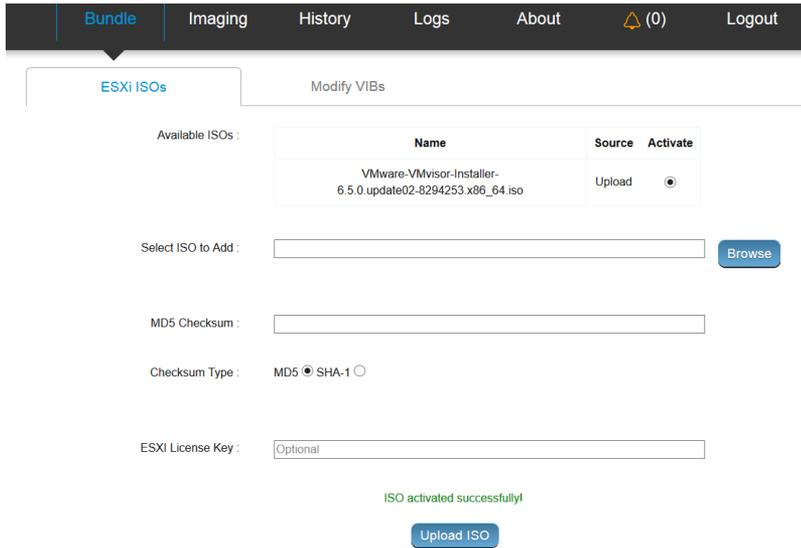
Procedure

- 1 In a web browser on the Windows machine that is connected to the Cloud Builder VM, navigate to `https://Cloud_Builder_VM_IP:8445/via`.

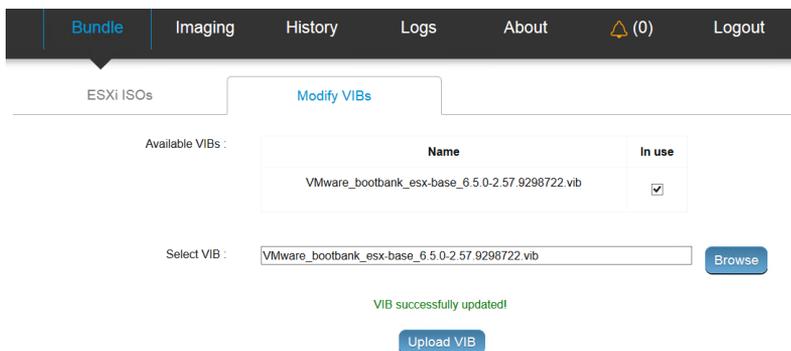
The VMware Imaging Appliance page displays.

- 2 Enter the admin credentials you provided when you deployed the Cloud Builder VM and click Log in.

- 3 Click **Bundle** and then click the **ESXi ISOs** tab.
- 4 Click **Browse** to locate and select the ISO.
- 5 Select the checksum type and enter the checksum.
- 6 Click **Upload ISO**.
- 7 When the uploaded ISO appears, select **Activate** to use the ISO for imaging servers.



- 8 Click the **Modify VIBs** tab.
- The steps for uploading VIBs are optional.
- 9 Click **Browse** to locate and select the VIB.
 - 10 Click **Upload VIB**.
 - 11 When the uploaded VIB appears, select **In use** to use the VIB for imaging servers.



What to do next

Use the selected ISO and VIB(s) to image servers for use with Cloud Foundation.

Image Servers with ESXi and VIBs

Once you have uploaded the required ESXi and VIB packages to the VMware Imaging Appliance, you can begin imaging servers. You can image an individual server, or multiple servers at the same time.

You can use VIA to image servers for use in the management domain and VI workload domains. The management domain requires a minimum of four servers. See the *VMware Cloud Foundation Planning and Preparation Guide* for more information about requirements.

Note When you image servers, VIA uses the ESXi ISO that you activated and the VIB(s) that you marked as **In use**.

Procedure

- 1 In a web browser on the Windows machine that is connected to the Cloud Builder VM, navigate to `https://Cloud_Builder_VM_IP:8445/via`.

The VMware Imaging Appliance page displays.

- 2 Enter the admin credentials you provided when you deployed the Cloud Builder VM and click Log in.
- 3 Click Imaging.
- 4 Enter the required information.

Name

Description

ESXI SERVER

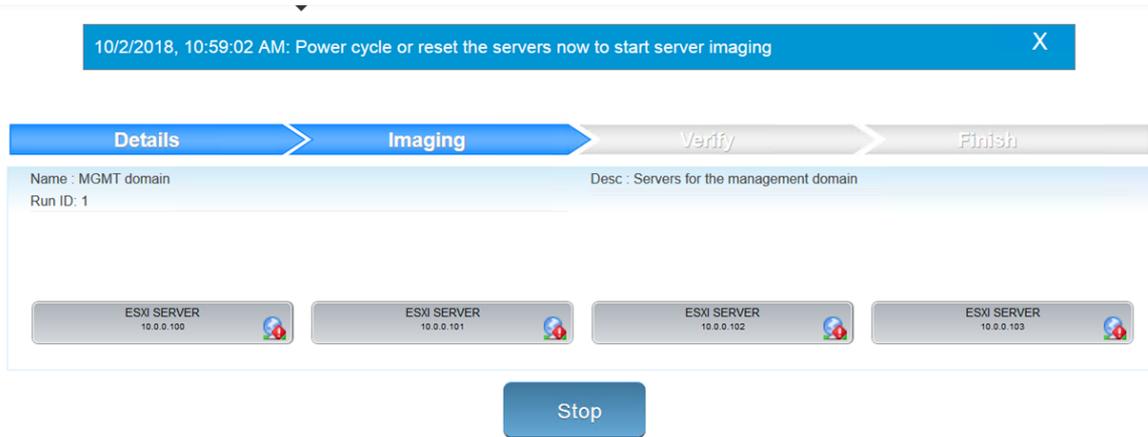
NTP Server: Number:

1	IP: <input type="text" value="10.0.0.100"/>	MAC: <input type="text" value="02:00:46:d2:08:73"/>	Hostname: <input type="text" value="esxi-1.vrack"/>	Host FQDN: <input type="text" value="esxi-1.vrack.vsphere.local"/>
2	IP: <input type="text" value="10.0.0.101"/>	MAC: <input type="text" value="02:00:46:f6:0c:9c"/>	Hostname: <input type="text" value="esxi-2.vrack"/>	Host FQDN: <input type="text" value="esxi-2.vrack.vsphere.local"/>
3	IP: <input type="text" value="10.0.0.102"/>	MAC: <input type="text" value="02:00:46:7b:c5:0f"/>	Hostname: <input type="text" value="esxi-3.vrack"/>	Host FQDN: <input type="text" value="esxi-3.vrack.vsphere.local"/>
4	IP: <input type="text" value="10.0.0.103"/>	MAC: <input type="text" value="02:00:46:2c:19:4a"/>	Hostname: <input type="text" value="esxi-4.vrack.vsphere"/>	Host FQDN: <input type="text" value="esxi-4.vrack.vsphere.local"/>

Option	Description
Name	Enter a name for the imaging job.
Number	Enter the number of servers you want to image with the selected ISO and VIBs.
Description	Enter a description for the imaging job.
NTP Server	Enter the IP address for the NTP server.
IP	Enter the IP address for the server.

Option	Description
MAC	Enter the MAC address for the server.
Hostname	Enter the hostname for the server.
Host FQDN	Enter the FQDN for the server.

- 5 Click **Start Imaging**.
- 6 When prompted, power cycle the server(s) to continue imaging.



VIA displays information about the progress of imaging. Click a server to view details. Once imaging is complete, VIA performs verification of the servers.

- 7 When verification is finished, click **Complete**.

What to do next

Perform post-imaging tasks.

Post-Imaging Tasks

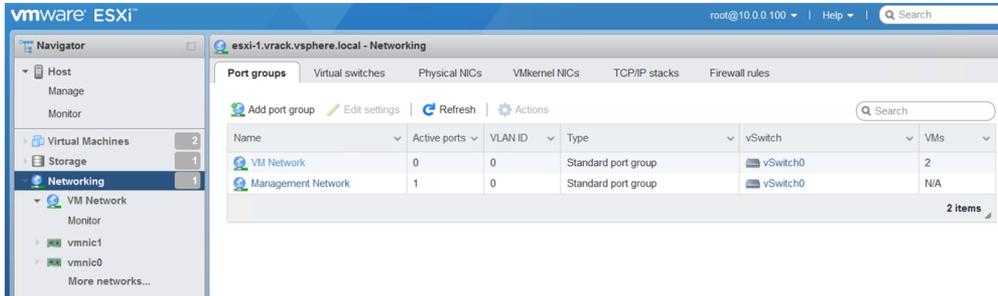
After you image your servers with ESXi and VIBs, you must perform some post-imaging tasks, depending on whether you use an untagged or a tagged management VLAN.

For imaging servers, the VMware Imaging Appliance requires an untagged VLAN. You can continue to use an untagged VLAN for management, or you can use a tagged VLAN.

Untagged Management VLAN

In this scenario, you use the same network for provisioning and management.

- Ensure that the Management Network and VM Network port groups on each host use the untagged VLAN (VLAN ID 0)



- Verify that your DNS and NTP server are routable to the management network and ESXi hosts can reach them. To configure a default gateway or static routes on your ESXi hosts, see <https://kb.vmware.com/kb/2001426>.

Tagged Management VLAN

In this scenario, you use an untagged VLAN for provisioning and a tagged VLAN for management.

- Modify the Management Network and VM Network port groups on each host to use the tagged VLAN
- Migrate the hosts from the provisioning network to the management network on the TOR switches
- Verify that your DNS and NTP server are routable to the management network and ESXi hosts can reach them. To configure a default gateway or static routes on your ESXi hosts, see <https://kb.vmware.com/kb/2001426>.

Host Management

7

To add hosts to the Cloud Foundation inventory, you must first create a network pool or expand the default network pool created during bring-up.

For information on network pools, see [About Network Pools](#).

You then commission hosts to Cloud Foundation. During the commissioning process, you associate hosts with a network pool. Commissioned hosts are added to the Cloud Foundation inventory. You can add these hosts to the management domain or to a VI workload domain. When a host is added to a workload domain, an IP address from the network pool's IP inclusion range is assigned to it.

This chapter includes the following topics:

- [About Network Pools](#)
- [Commission Hosts](#)
- [Decommission Hosts](#)
- [Cleaning Up Decommissioned Hosts](#)
- [View Host Inventory](#)

About Network Pools

Network pools automatically assign static IP addresses to vMotion, vSAN, and NFS vmkernel ports so that you don't need to enter IP addresses manually when creating a VI workload domain or adding a host or cluster to a workload domain.

A network pool is a collection of a set of subnets within an L2 domain. Depending on the storage option you are using, it includes information about subnets reserved for the vMotion and vSAN or NFS networks that are required for adding a host to the Cloud Foundation inventory.

Table 7-1. Information Required for a Network Pool

Storage Being Used	Required Networks in Network Pool
vSAN	vMotion and vSAN
NFS	vMotion and NFS

Table 7-1. Information Required for a Network Pool (continued)

Storage Being Used	Required Networks in Network Pool
vSAN and NFS	vMotion, vSAN, and NFS
VMFS on FC	vMotion only or vMotion and NFS

The network pool also contains a range of IP addresses, called an inclusion range. IP addresses from the inclusion ranges are assigned to the vMotion and vSAN or NFS vmkernel ports on the host. The use of inclusion ranges allows you to limit the IP addresses that will be consumed from a given subnet. You can add more inclusion ranges in order to expand the use of the provided subnet.

A default network pool (named bringup-networkpool) is created during bring-up. This network pool is automatically associated with the management domain. Network information for this network pool is based on the deployment parameter sheet you provided during bring-up. This network pool contains vMotion and vSAN networks only - an NFS network is not supported in this network pool. If you have a single L2 domain in your environment for management workload domain vSAN and vMotion networks or if you want to expand the management domain by adding a host, you can expand this default network pool.

In order to create a workload domain with hosts in a different L2 domain than the management domain, or if you want to use external NFS or VMFS on FC storage, you must create a new network pool. A network pool can contain both vSAN and NFS networks.

For NSX for vSphere workload domains, all hosts in a cluster must be associated with the same network pool. For NSX-T workload domains, you can use the Cloud Foundation API to select hosts from different network pools, as long as those network pools have the same VLAN ID and MTU settings.

All hosts in a cluster must be associated with the same network pool. However, a workload domain can contain multiple clusters, each with its own network pool. You may want to have multiple clusters within a workload domain to provide separate fail over domains (i.e. a VM only fails over between hosts in a cluster). Multiple clusters also provide isolation for security reasons and are also useful for grouping servers of a particular type of configuration together. Multiple clusters can also be used to handle growth. Original servers used in the first cluster may get outdated at some point. Newer server models can then be added in a new cluster to the workload domain and workloads can be migrated at a leisurely pace.

Sizing a Network Pool

Properly sizing a network pool is critical to prevent future issues in the environment due to insufficient IP addresses. Care must be taken when defining the subnets for a network pool as the subnet cannot be changed after it is deployed. The scope of IP addresses used from the defined subnet can be limited by the definition of one or more inclusion ranges. Thus, it is recommended that you begin with defining a larger subnet than what is initially required and utilize the inclusion ranges to limit use. This will provide you the capability to grow with demand as needed.

You begin sizing a network pool by determining the number of hosts that you will have in each cluster. A workload domain must contain a minimum of one cluster. As each cluster leverages vSAN for storage, the minimum number of hosts within a cluster is three. The exception to this rule is the management workload domain. It is recommended that the management workload domain contain a minimum of four hosts. This allows for an additional level of availability for the critical infrastructure components. A cluster can be expanded to the maximum number of hosts supported by vCenter, which is currently 64 hosts.

Allocate a minimum of one IP address per host plus enough additional IP addresses to account for growth and expansion of the environment. Ensure that the subnet defined provides enough unused IP addresses and that appropriate inclusion ranges are defined. Note that some of the IP addresses within the subnet will be used for other purposes, such as defining the gateway address, firewalls, or other entities. Use care not to conflict with these addresses.

Here are some important considerations for determining the size of your network pool:

- Type of network architecture
- Physical switch details
 - Are they managed or non-managed
 - Do they support L3 (this may require a license)
 - Number of ports
- Where the network switches are placed (at the top of the rack or at the end of a row)
- Where the default gateway is created
- Number of hosts that can be placed in each rack or L2 domain
- Number of hosts required in a cluster
- Whether the network switches will be shared with non-Cloud Foundation hosts
- Number of workload domains you plan on creating

Create a Network Pool

A network pool must include vMotion network information. Depending on the type of storage you are using, you may also need to provide network information for vSAN and NFS.

The subnet in a network pool cannot overlap the subnet of another pool.

Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Network Settings..**
- 2 Click **Create Network Pool.**
- 3 Enter a name for the network pool.

4 Select the network type(s).

You can include both vSAN and NFS network information in the same network pool, or create separate network pools for vSAN and NFS. For VMFS on FC storage, select vMotion only or vMotion and NFS.

5 Provide the following network information for the selected network type(s).

- a Enter a VLAN ID between 1 and 4094.
- b Enter an MTU between 1500 and 9216.
- c In the **Network** field, enter a subnet IP address.
- d Enter the subnet mask.
- e Enter the default gateway.
- f Enter an IP address range from which an IP address can be assigned to hosts that are associated with this network pool.

The IP address range must be from within the specified subnet. You cannot include the IP address of the default gateway in the IP address range. You can enter multiple IP address ranges.

Note Ensure that you have entered the correct IP address range. IP ranges cannot be edited after the network pool is created.

6 Click **Save**.

View Network Pool Details

You can view vSAN and vMotion network details for a network pool as well as the total number of used and available IP addresses.

Procedure

- 1** On the SDDC Manager Dashboard, click **Administration > Network Settings**.

- 2** Click the arrow to the left of the pool name.

A high-level summary of the network pool's vSAN and vMotion network information is displayed.

- 3** Click the name of a network pool.

Network pool details are displayed.

Edit a Network Pool

You can add an IP inclusion range to a network pool. No other parameters can be modified.

Procedure

- 1** On the SDDC Manager Dashboard, click **Administration > Network Settings**.

- 2 Hover your mouse in the network pool row that you want to edit.

A set of three dots appear on the left of the pool name. Click these dots and then click **Edit**.

- 3 Enter an IP inclusion range and click **Add**.
- 4 Click **Save**.

Delete a Network Pool

You can delete a network pool if none of the hosts with an IP address from this pool belong to a workload domain. The default pool created during bring-up cannot be deleted.

Prerequisites

Ensure that the hosts in the network pool are not assigned to a workload domain. To verify this, navigate to **Administration > Network Settings** and confirm that the **Used IPs** for the network pool is 0.

Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Network Settings**.
- 2 Hover your mouse in the network pool row that you want to delete.
A set of three dots appear on the left of the pool name. Click these dots and then click **Delete**.

Commission Hosts

Adding hosts to the Cloud Foundation inventory is called commissioning. You can add hosts individually, or use a JSON template to add multiple hosts at once. You can commission a maximum of 32 hosts at a time.

The hosts that you want to commission must meet a set of criteria. After you specify host details and select the network pool to associate a host with, Cloud Foundation validates and commissions each host. Each host is added to the free pool and is available for workload domain creation.

The storage type you select for a host (vSAN, NFS, VMFS on FC), must be supported by its associated network pool. A network pool can support both vSAN and NFS. For VMFS on FC storage, the network pool must be vMotion only or vMotion and NFS.

- Hosts that use vSAN storage can only be used with vSAN-based workload domains.
- Hosts that use NFS storage can only be used with NFS-based workload domains.
- Hosts that use VMFS on FC storage can only be used with VMFS on FC-based workload domains.

Note The management domain can only include hosts that use vSAN storage.

Prerequisites

Ensure that each host you are commissioning meets the following criteria.

- Hosts for vSAN-based workload domains are vSAN-compliant and certified on the VMware Hardware Compatibility Guide.
- Hosts for NFS-based workload domains are certified on the VMware Hardware Compatibility Guide.
- Hosts for VMFS on FC-based workload domains are certified on the VMware Compatibility Guide. In addition, the hosts must have supported FC cards (Host Bus Adapters) and drivers installed and configured. For compatible FC cards, see the VMware Compatibility Guide.
- Host has the drivers and firmware versions specified in the VMware Hardware Compatibility Guide.
- Hardware health status is healthy without any errors.
- A supported version of ESXi is installed on the host. See the *VMware Cloud Foundation Release Notes* for information about supported versions.
- Two NIC ports with a minimum 10 Gbps speed. One port must be free and the other port must be configured on a standard switch. This switch should be restricted to the management portgroup.

Starting with Cloud Foundation 3.9.1, you can commission hosts with more than two NICs using APIs. For information on multi-pNIC support, see [Separating Traffic by Using Multiple vDSes](#).

- Management IP address is configured on the first NIC port.
- Host is configured with appropriate gateway. The gateway must be part of the management subnet.
- SSH and syslog are enabled.
- DNS is configured for forward and reverse lookup and FQDN.
- All disk partitions on HDD and SSD are deleted.

Note You must have a network pool available in order to commission a host.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Hosts**.
- 2 Click **Commission Hosts**.
- 3 Confirm that hosts to be commissioned meet each criterion in the checklist and select the check boxes.
- 4 Click **Proceed**.

- 5 Select whether you want to add hosts one at a time, or import a JSON file to add multiple hosts at once.

Option	Description
Add new	<p>Manually enter the following information for the host you want to add:</p> <ul style="list-style-type: none"> ■ FQDN ■ Network pool (choose an existing network pool from the list) ■ User name and password (root credentials) ■ Storage type (vSAN, NFS, or VMFS on FC) <p>Host Addition and Validation 🔗</p> <hr/> <p>▼ Add Hosts</p> <p>You can either choose to add host one at a time or download JSON template and perform bulk commission.</p> <p><input checked="" type="radio"/> Add new <input type="radio"/> Import</p> <p>Host FQDN: <input type="text" value="esxi-5.vrack.vsphere.local"/></p> <p>Storage Type: <input checked="" type="radio"/> vSAN <input type="radio"/> NFS <input type="radio"/> VMFS on FC</p> <p>Network Pool Name ^①: <input type="text" value="bringup-networkpool"/></p> <p>User Name: <input type="text" value="root"/></p> <p>Password: <input type="password" value="....."/></p> <p style="text-align: right;">ADD</p>

Click **Add**.

You can now add more hosts or proceed to the next step.

Import

- a Click the link to download the JSON template.
- b Open the JSON template file and enter information about the hosts to add.
 - FQDN
 - User name and password (root credentials)
 - Storage type (vSAN, NFS, or FC)
 - Network pool name

```

1  {
2      "hostsSpec": [
3          {
4              "hostfqdn": "esxi-5.vrack.vsphere.local",
5              "username": "root",
6              "password": "Er5!x98b",
7              "storageType": "vSAN",
8              "networkPoolName": "bringup-networkpool"
9          },
10         {
11             "hostfqdn": "esxi-6.vrack.vsphere.local",
12             "username": "root",
13             "password": "B89x!5rE",
14             "storageType": "NFS",
15             "networkPoolName": "sfo-networkpool"
16         }
17     ]
18 }
    
```

Option	Description
	c Click Browse to locate and select the JSON file containing host information.
	d Click Upload .

The host or hosts appear in the **Hosts Added** section.

- Verify that the server fingerprint is correct for each host and then click the confirm fingerprint icon .

- Click **Validate All**.

Cloud Foundation validates the host information you provided. Each host is marked as **Valid** or **Invalid**.

For invalid hosts, you can correct the problem and validate again, or select the host and click **Remove** to proceed with commissioning the valid hosts.

- Click **Next** to review the host information and then click **Commission** to begin commissioning.

The Hosts page appears and the status of the commission task is displayed. Click **View Status in Task** to display the task bar.

Results

The commissioned hosts are added to the host table. The host belongs to a free pool until you assign it to a workload domain.

Decommission Hosts

Removing hosts from the Cloud Foundation inventory is called decommissioning. You can decommission a host for maintenance work or if you want to add it to another network pool. If you want to re-use a host in a different workload domain, you must decommission the host and clean it up before adding it to the workload domain.

Prerequisites

The hosts that you want to decommission must not be assigned to a workload domain. If a host is assigned to a workload domain, you must remove it before you can decommission it. See [Remove a Host from a Cluster in a Workload Domain](#).

Procedure

- On the SDDC Manager Dashboard, click **Inventory > Hosts**.
- Click **Unassigned Hosts**.
- In the hosts table, select the host(s) you want to decommission.
- Click **Decommission Selected Hosts**.

5 Click **Confirm**.

The Hosts page appears and the status of the decommission task is displayed. Click **View Status in Task** to display the task bar.

What to do next

Clean the decommissioned host before adding it to a workload domain. See [Cleaning Up Decommissioned Hosts](#).

Cleaning Up Decommissioned Hosts

Before you can add a decommissioned host to a workload domain, you must clean it up.

The best way to clean up a decommissioned host is by using the SoS utility on the SDDC Manager VM. If the SoS utility is unable to clean up the host for some reason, you can use the Direct Console User Interface (DCUI) on the host to perform the cleanup.

Clean up a Decommissioned Host Using the SoS Utility

A decommissioned host must be cleaned up before it can be assigned to a workload domain. You can use the SoS utility on the SDDC Manager VM to perform host cleanup.

Prerequisites

Gather the following information for each host that you want to clean up:

- IP address
- User name and password (root credentials)

Procedure

- 1 Using SSH, log in to the SDDC Manager VM with the following credentials:
Username: **vcf**
Password: use the password specified in the deployment parameter sheet
- 2 Enter **su** to switch to the root user.
- 3 Change to the `/opt/vmware/sddc-support` directory.
- 4 Edit `decommissioned_host_cleanup_sample.json` to include information for the host(s) you want to clean up.
- 5 Type the following command:

```
./sos --cleanup-decommissioned-host  
/opt/vmware/sddc-support/decommissioned_host_cleanup_sample.json
```

What to do next

You can now commission the host to the Cloud Foundation inventory and add it to a workload domain.

Clean up a Decommissioned Host Using the Direct Console User Interface

A decommissioned host must be cleaned up before it can be assigned to a workload domain. You can use the Direct Console User Interface (DCUI) on the host to perform host cleanup.

Prerequisites

- You must have access to Direct Console User Interface (DCUI) on the host.
- Gather the following information for the decommissioned host:
 - IP address
 - root password
 - network configuration - netmask, gateway, and DNS
 - VLAN ID

Procedure

- 1 Log in to the DCUI.
- 2 Navigate to the Troubleshooting Options page and enable ESXI shell.
- 3 Press Alt-F1 to get to the prompt to run command line steps.
- 4 Clean up vSAN with the following command.

```
#vdq -i
#esxcli vsan storage remove -s SSD Device Name
```

For example:

```
[root@esx-6:/tmp] vdq -i
[
  {
    "SSD" : "naa.55cd2e414dc36b15",
    "MD" : [
      "naa.55cd2e414d7abb5d",
      "naa.55cd2e414d7aa215",
      "naa.55cd2e414d7abb46",
    ]
  },
  {
    "SSD" : "naa.55cd2e414dc36d53",
    "MD" : [
      "naa.55cd2e414d705c35",
      "naa.55cd2e414d7aa1eb",
      "naa.55cd2e414d7abb10",
    ]
  }
]
```

```

    ],
  ],
[root@esx-6:/tmp] esxcli vsan storage remove -s naa.55cd2e414dc36b15
[root@esx-6:/tmp] esxcli vsan storage remove -s naa.55cd2e414dc36d53
[root@esx-6:/tmp] vdq -i
[
]

```

- 5 Reset the system configuration and the root password by running the commands below.


```
/bin/firmwareConfig.sh --reset
```

When you reset the configuration, the software overrides all your network configuration changes, deletes the password for the administrator account (root), and reboots the host.
- 6 Press Alt-F2 to return to the DCUI.
- 7 Reset the root password. This password was deleted during step 5.
- 8 Configure the following network details to the same values that were set on the host before the factory reset.
 - VLAN
 - Set static IPv4 address
 - IP address
 - netmask
 - gateway
- 9 Apply the changes.
- 10 Restart the management network by selecting the Restart Management Network option on the main DCUI page.
- 11 On the Troubleshooting Options page, enable SSH on the host.
- 12 Disable ESXi shell.

What to do next

You can now commission the host to the Cloud Foundation inventory and add it to a workload domain.

View Host Inventory

The Hosts page displays details about all the hosts in your Cloud Foundation system, including CPU utilization and memory usage across all hosts, as well as the total number of hosts used and unallocated.

For each host, the Hosts page displays the following information:

- FQDN name

- IP address
- The network pool to which the host belongs
- Current status
- Host state, or the workload domain to which it is allocated
- Cluster or more specifically, the domain cluster to which it is assigned
- Host-specific CPU and memory usage
- Storage type

The Hosts page also provides controls for commissioning hosts.

Procedure

- 1 From the the SDDC Manager Dashboard, navigate to **Inventory > Hosts**.

The Hosts page appears.

- 2 Navigate directly to pages related to a specific host.

For example:

- To jump to the details page for the domain to which a listed host belongs, click the domain name link in the Host State column. For information about viewing workload domains, see [View Workload Domain Details](#).
- To jump to the details page for the domain cluster to which a listed host belongs, click the cluster name in the Cluster column. For information about clusters, see [Expand a Workload Domain](#).
- To quickly view network assignment details for a specific host, click the info icon next to the value in the Network Pool column.

- 3 To view the details of a specific host, click the FQDN name in the list.

The host details page appears, displaying the following information:

- A chart showing total and used CPU capacity.
- A chart showing total and used memory capacity.
- A summary of the networks (vSAN, vMotion, and Management) to which the host belongs and its IP address on those networks.
- The manufacturer and model of the host.
- Storage information including capacity and cache tiers.

Note Below the page title, the host details page also provides quick links to the network pool and the workload domain cluster to which the host belongs.

- 4 (Optional) To decommission the host from the host details page, click **Actions** near the page name and select **Decommission**.

For details, see [Decommission Hosts](#).

- 5 (Optional) To view host VM details, click **Actions** near the page name and select **Open in ESXi Client**.

The ESXi Client opens.

Working with the Management Domain and VI Workload Domains



The management domain and deployed VI workload domains are logical units that carve up the compute, network, and storage resources of the Cloud Foundation system. The logical units are groups of ESXi hosts managed by vCenter Server instances with specific characteristics for redundancy and VMware SDDC best practices.

For information on Horizon domains, see [Chapter 10 Working with Horizon Domains](#).

For information on deploying an Enterprise PKS solution, see [Chapter 9 Working with VMware Enterprise PKS](#).

The management domain is created by default during bring-up. The Cloud Foundation software stack is deployed on the management domain. Additional infrastructure virtual machines which provide common services, such as backup or security appliances, can be deployed in the management domain as well.

The management domain and workload domains include these VMware capabilities by default:

VMware vSphere[®] High Availability (HA)

This feature supports distributed availability services for a group of ESXi hosts to provide rapid recovery from outages and cost-effective high availability for applications running in virtual machines. Out of the box, Cloud Foundation provides a highly available environment for workload domains. There may be additional settings (not set by default) that can increase availability even further. For more information about vSphere HA, see the *vSphere Availability* documentation at <https://docs.vmware.com/en/VMware-vSphere/>.

VMware vSphere[®] Distributed Resource Scheduler™ (DRS)

This feature dynamically allocates and balances computing capacity across a group of hardware resources aggregated into logical resource pools or clusters. Clusters are the primary unit of operation in Cloud Foundation. DRS continuously monitors use across resource pools and allocates available resources among the virtual machines based on predefined rules that reflect business needs and changing priorities. When a virtual machine experiences an increased load, vSphere DRS automatically allocates additional resources by redistributing virtual machines among the physical servers in the resource pool. For more

information about DRS, see the *vSphere Resource Management* documentation at <https://docs.vmware.com/en/VMware-vSphere/>.

VMware vSAN[®]

This component aggregates local storage disks in a group of ESXi hosts to create a storage pool shared across all hosts in that group. For more information about vSAN, see the *VMware vSAN* documentation at <https://docs.vmware.com/en/VMware-vSAN/>.

VMware NSX Automated Backup

It is crucial to take backups of all NSX components to restore the system to its working state in the event of a failure. VMware Cloud Foundation automatically configures the NSX Manager applications to back up their state to a user-supplied SFTP server. These backups contain all of the NSX configuration, including controllers, logical switching and routing entities, security, firewall rules, and everything else that you configure within the NSX Manager UI or API. The vCenter database and related elements like the virtual switches need to be backed up separately.

Each Cloud Foundation instance is one SSO domain to which all vCenter Servers are joined. The maximum number of supported workload domains and vCenter Servers per Cloud Foundation instance depends on the vSphere version in the management cluster. For more information, see the *Configuration Maximums vSphere* document.

Note if you use cross vCenter vMotion between two VI workload domains with dissimilar hardware, you must enable EVC on the corresponding clusters. See [Enable EVC on an Existing Cluster](#) in the vSphere product documentation. You can enable EVC on the management domain by selecting the appropriate value in the Deploy Parameters tab of the Deployment Parameter spreadsheet. For more information, see the *VMware Cloud Foundation Architecture and Deployment Guide*.

Procedure

1 Adding Virtual Machines to the Management Domain

You can add virtual machines to the management domain as desired. Commonly, these virtual machines provide infrastructure services such as backup or security throughout the solution. To prevent resource conflicts between the core Cloud Foundation services, these additional virtual machines are added to the Compute-ResourcePool. This resource pool is automatically created during bring-up for this purpose.

2 About VI Workload Domains

In the VI Configuration wizard, you specify the storage, name, compute, and NSX platform details for the VI workload domain. Based on the selected storage, you provide vSAN parameters, NFS share details, or VMFS on FC datastore information. You then select the hosts and licenses for the workload domain and start the creation workflow.

3 Deploy and Configure NSX Edges in Cloud Foundation

For an NSX-T VI workload domain, NSX Edges are required to enable overlay VI and public networks for north-south traffic. A minimum of two NSX Edge nodes are required for high availability and redundancy.

4 View Workload Domain Details

The Workload Domains page displays high level information about the workload domains in the Cloud Foundation system. CPU, memory, and storage utilized by the workload domain is also displayed here.

5 Delete a VI Workload Domain

When you delete a workload domain, the clusters within that workload domain are deleted and the hosts are returned to the free pool.

6 View Cluster Details

The cluster page displays high level information about the cluster as well as the hosts that form that cluster. CPU, memory, and storage utilization for this cluster is also displayed here.

7 Shrink a Workload Domain

You can reduce the management domain or a VI workload domain by removing a host from a vSphere cluster in the workload domain or by deleting a vSphere cluster.

8 Expand a Workload Domain

You can expand the management domain or a VI workload domain to add resources to support additional workloads or availability.

Adding Virtual Machines to the Management Domain

You can add virtual machines to the management domain as desired. Commonly, these virtual machines provide infrastructure services such as backup or security throughout the solution. To prevent resource conflicts between the core Cloud Foundation services, these additional virtual machines are added to the Compute-ResourcePool. This resource pool is automatically created during bring-up for this purpose.

Note You must be careful when adding virtual machines to the management domain. You do not want to consume excessive resources that would obstruct standard operations. Excess capacity consumption can cause failures of virtual machine fail overs in the event of a host failure or maintenance action.

You can add capacity to the management domain by adding a host(s) in order to expand the management workload domain. To expand the management domain, see [Expand a Workload Domain](#).

Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Inventory > Workload Domains**.
- 2 In the workload domains table, click **MGMT**.

3 On the MGMT page, click the **Services** tab.

4 Click the vCenter link.

This opens the vSphere Web Client for the management domain.

5 Create a VM.

See *Create a New Virtual Machine* in *vSphere Resource Management*.

Note Do not move any of the Cloud Foundation management VMs into the resource pool.

6 Move the VM to the resource pool.

See *Add a Virtual Machine to a Resource Pool* in *vSphere Resource Management*.

Note Do not move any of the Cloud Foundation management VMs to the newly created resource pool.

About VI Workload Domains

In the VI Configuration wizard, you specify the storage, name, compute, and NSX platform details for the VI workload domain. Based on the selected storage, you provide vSAN parameters, NFS share details, or VMFS on FC datastore information. You then select the hosts and licenses for the workload domain and start the creation workflow.

The workflow automatically:

- Deploys an additional vCenter Server Appliance for the new workload domain within the management domain.

By using a separate vCenter Server instance per workload domain, software updates can be applied without impacting other workload domains. It also allows for each workload domain to have additional isolation as needed.

- Connects the specified ESXi servers to this vCenter Server instance and groups them into a cluster. Each host is configured with the port groups applicable for the workload domain.
- Configures networking on each host.
- Configures vSAN, NFS, or VMFS on FC storage on the ESXi hosts.
- For each NSX for vSphere workload domain, the workflow deploys an NSX Manager in the management domain and three NSX controllers on the ESXi datastore. The workflow also configures an anti-affinity rule between the controller VMs to prevent them from being on the same host for High Availability.
- After deploying NSX Manager, the workflow configures the NSX Manager to back up its state periodically to an SFTP server. In the earlier releases, these backups were written to the SFTP server built into SDDC Manager. From Release 3.9 onwards, you can register an external SFTP server with SDDC Manager for these backups. If you register an external SFTPI server,

SDDC Manager uses it with all existing NSX Managers with it and will use this external SFTP server when deploying additional NSX Managers. The built-in SFTP server is used until you configure an external one. For more information about the SFTP server feature, see [Configure an External SFTP Server for NSX Manager Backups](#).

- For the first NSX-T VI workload domain, the workflow deploys a cluster of three NSX-T Managers in the management domain. The workflow also configures an anti-affinity rule between the NSX Manager VMs to prevent them from being on the same host for High Availability. All subsequent NSX-T workload domains share this NSX-T Manager cluster.

For an NSX-T workload domain, NSX Edges are required to enable overlay VI networks and public networks for north-south traffic. NSX Edges are not deployed automatically for an NSX-T VI workload domain. You can deploy them manually after the VI workload domain is created. Subsequent NSX-T VI workload domains share the NSX-T Edges deployed for the first workload domain.

- Licenses and integrates the deployed components with the appropriate pieces in the Cloud Foundation software stack.

The result is a workload-ready SDDC environment.

Note You can only perform one workload domain operation at a time. For example, while creating a new workload domain, you cannot add a cluster to any other workload domain.

Procedure

1 [Prerequisites for a Workload Domain](#)

This section lists pre-requisites for a VI workload domain.

2 [Additional Prerequisites for an NSX-T Based Workload Domain](#)

You must download the NSX-T binaries before creating the VI workload domain. The procedure you follow depends on the Cloud Foundation version in your environment.

3 [Start the VI Configuration Wizard](#)

Start the VI Configuration wizard and select the storage type for the workload domain.

Prerequisites for a Workload Domain

This section lists pre-requisites for a VI workload domain.

NSX-T VI workload domains have additional pre-requisites. See [Additional Prerequisites for an NSX-T Based Workload Domain](#).

- Do one of the following.
 - For NSX-V workload domains, a DHCP server scope must be configured for every NSX-V VXLAN Tunnel End Point (VTEP) VLAN within the VI workload domain. When NSX-V creates VXLAN VTEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.

- For NSX-T workload domains, a DHCP server scope must be configured for every NSX-T Tunnel End Point (TEP) VLAN within the VI workload domain. NSX-T uses the Geneve encapsulation protocol and each TEP interfaces requires an IP address assigned from a reliable DHCP server.
- A minimum of three hosts marked with the appropriate storage must be available in your Cloud Foundation inventory. To create a VI workload domain with NFS storage, the hosts must be commissioned with NFS as the storage type and must be associated with an NFS network pool. To create a VI workload domain with vSAN storage, the hosts must be commissioned with vSAN as the storage type and must be associated with an vSAN network pool. To create a VI workload domain with VMFS on FC storage, the hosts must be commissioned with VMFS on FC as the storage type and must be associated with a vMotion only or vMotion and NFS network pool. For information on adding hosts to your inventory, see [Chapter 7 Host Management](#).
- There must be a free uplink on each host to be used for the workload domain.
- If the management domain in your environment has been upgraded to a version different from the original installed version, you must download a VI workload domain install bundle before you can create a VI workload domain. See [Chapter 13 Downloading an Install Bundle](#).
- Decide on a name for your VI workload domain. Each VI workload domain must have a unique name. It is good practice to include region and site information in the name because resource object names (such as host and vCenter names) are generated based on the VI workload domain name. The name can be three to 20 characters long and can contain any combination of the following:
 - Lowercase alphabetic characters
 - Uppercase alphabetic characters
 - Numbers

Note Spaces are not allowed in any of the names you specify when creating a VI workload domain.

- Decide on the following passwords - refer to the appropriate table for the Cloud Foundation version in your environment.
 - vCenter root password
 - NSX Manager admin password
 - NSX Manager password to enable administrator privileges for NSX Manager (only for NSX-V)

Table 8-1. Passwords for Cloud Foundation

Account	Password Requirements
vCenter root	<ol style="list-style-type: none"> 1 Length 8-20 characters 2 Must include: <ul style="list-style-type: none"> ■ mix of upper-case and lower-case letters ■ a number ■ a special character
NSX-V Manager admin and	<ol style="list-style-type: none"> 1 Length 8-12 characters 2 Must include: <ul style="list-style-type: none"> ■ mix of upper-case and lower-case letters ■ a number ■ a special character ■ exclude_char such as { } [] () / \ ' " ` ~ , ; : . < >
NSX-T Manager	<ol style="list-style-type: none"> 1 Minimum length 12 characters 2 Must include: <ul style="list-style-type: none"> ■ at least one lowercase and one uppercase letter ■ a number ■ a special character ■ exclude_char such as { } [] () / \ ' " ` ~ , ; : . < > ■ at least five different characters 3 Must not include: <ul style="list-style-type: none"> ■ a dictionary word ■ a palindrome ■ more than four monotonic character sequences

- Based on the Cloud Foundation version in your environment, gather the information that you need for the workload domain creation workflow.

vCenter IP address, DNS name, subnet mask, and default gateway

Three NSX Managers IP address, DNS name, subnet mask, and default gateway

NSX Manager Virtual IP (VIP) address

- The IP addresses and Fully Qualified Domain Names (FQDN) for the vCenter and NSX Manager instances to be deployed for the VI Workload domain must be resolvable by DNS.
- If you are using NFS storage for the workload domain, you need the following information:
 - Datastore name
 - Path to the NFS share
 - IP address of the NFS server

The NFS share and server must be accessible from the Cloud Foundation network. You must have read/write permission to the NFS share because NSX controllers will be deployed there.

- If you are using VMFS on FC storage for the workload domain, you need the datastore name.

- You must have valid license keys for the following products:
 - vCenter Server
 - NSX for vSphere or NSX-T
 - vSAN (No license required for NFS or VMFS on FC)
 - vSphere

Because vSphere and vSAN licenses are per CPU, ensure that you have sufficient licenses for the ESXi hosts to be used for the workload domain. See [Chapter 5 License Management](#).

- (Optional) Enable vRealize Log Insight logging for workload domains. See [Enable vRealize Log Insight Logging for Workload Domains](#).
- If you have upgraded the management domain in your environment to a later release, download the VI workload domain install bundle to deploy later versions of the software components instead of the versions in your original Cloud Foundation installation. See "Download Bundles" in the *VMware Cloud Foundation Upgrade Guide*.

Additional Prerequisites for an NSX-T Based Workload Domain

You must download the NSX-T binaries before creating the VI workload domain. The procedure you follow depends on the Cloud Foundation version in your environment.

Procedure

- 1 Ensure that you have a separate VLAN dedicated for Edge TEPs and have routing configured with MTU 9000 between host VTEPs and Edge VTEP VLANs.
- 2 Download the NSX-T Manager 2.5 install bundle. See [Download Bundles](#) in the *VMware Cloud Foundation Upgrade Guide*.

Although Cloud Foundation supports NSX-T 2.4.2 for existing workload domains, you must have NSX-T 2.5 to deploy a new NSX-T based workload domain.

- 3 Add an NSX-T license key to SDDC Manager. See [Add License Keys for the Software in Your Cloud Foundation System](#).

Physical Networking Design for NSX-T Workload Domains

This section describes best practises for designing the physical network for NSX-T workload domains and Enterprise PKS.

Top of Rack Physical Switches

Consider the following best practises when configuring the top of rack (ToR) switches.

- Configure redundant physical switches to enhance availability.
- Configure switch ports that connect to ESXi hosts manually as trunk ports.

- Modify the Spanning Tree Protocol (STP) on any port that is connected to an ESXi NIC to reduce the time to transition ports over to the forwarding state, for example using the Trunk PortFast feature in Cisco physical switches.
- Provide DHCP or DHCP Helper capabilities on all VLANs used by TEP VMkernel ports. This setup simplifies the configuration by using DHCP to assign IP address based on the IP subnet in use.
- Configure jumbo frames on all switch ports, inter-switch link (ISL), and switched virtual interfaces (SVIs).

Top of Rack Connectivity and Network Settings

Each ESXi host is connected redundantly to the ToR switches SDDC network fabric by two 25 GbE ports. Configure the ToR switches to provide all necessary VLANs using an 802.1Q trunk. These redundant connections use features in vSphere Distributed Switch and NSX-T to guarantee that no physical interface is overrun and available redundant paths are used.

VLANs and Subnets

Each ESXi host uses VLANs and corresponding subnets. Follow these guidelines for VLANs and subnets.

- Use only /24 subnets to reduce confusion and mistakes when handling IPv4 subnetting.
- Use the IP address .254 as the (floating) interface with .252 and .253 for Virtual Router Redundancy Protocol (VRRP) or Hot Standby Routing Protocol (HSRP).
- Use the RFC1918 IPv4 address space for these subnets and allocate one octet by region and another octet by function.

Access Port Network Settings

Configure additional network settings listed in the table below on the access ports that connect the ToR switches to the corresponding servers.

Table 8-3. Additional Network Settings to be Configured on VLANs and Subnets

Setting	Description
Spanning Tree Protocol (STP)	Although this design does not use the Spanning Tree Protocol, switches usually include STP configured by default. Designate the access ports as trunk PortFast.
Trunking	Configure the VLANs as members of a 802.1Q trunk with the management VLAN acting as the native VLAN.
MTU	Set MTU for all VLANs and SVIs (Management, vMotion, Geneve, and Storage) to jumbo frames for consistency.
DHCP Helper	Configure a DHCP helper (DHCP relay) on all TEP VLANs.

Routing Protocols

NSX-t supports BGP only. Configure BGP per the guidelines below. The Layer 3 device (for example, the ToR switch) must support BGP.

DHCP

Set the DHCP helper (relay) to point to a DHCP server by IPv4 address.

Physical and Logical Networking

- Implement the following physical network architecture:
 - One 25 GbE (10 GbE minimum) port on each ToR switch for ESXi host uplinks.
 - No EtherChannel (LAG/LACP/vPC) configuration for ESXi host uplink.
 - Use two ToR switches for each rack for redundancy.
- Implement the following logical network architecture:
 - Use VLANs to segment physical network functions.

Static IP Addresses, DNS records, and NTP time source

- Use a physical network that is configured for BGP routing adjacency.
- Use two ToR switches for each rack.
- Use VLANs to segment physical network functions.

Jumbo Frames

You must configure jumbo frames end-to-end. Select an MTU that matches the MTU of the physical switch ports. Note that the Geneve overlay requires an MTU value of 1600 bytes or greater.

Virtual Infrastructure

The default cluster of the NSX-T workload domain contains the NSX-T Edge nodes. If you deploy an Enterprise PKS solution on an NSX-T workload domain, the Enterprise PKS workload domains are also in the default cluster.

Start the VI Configuration Wizard

Start the VI Configuration wizard and select the storage type for the workload domain.

Procedure

- 1 On the SDDC Manager Dashboard, click **+ Workload Domain** and then click **VI Virtual Infrastructure**.
- 2 Select the storage type and click **Begin**.

Specify Name

Provide a name for the VI workload domain, cluster, and organization.

Prerequisites

Verify that you have met the prerequisites described in [About VI Workload Domains](#).

Procedure

- 1 Type a name for the VI workload domain, such as **sfo01**. The name must contain between 3 and 20 characters.

It is good practise to include location information in the name since resource object names (such as host and vCenter names) are generated on the basis of the VI workload domain name.

- 2 Type a name for the VI cluster. The name must contain between 3 and 20 characters.
- 3 (Optional) Type a name for the organization that requested or will use the virtual infrastructure, such as **Finance**. The name must contain between 3 and 20 characters.
- 4 Click **Next**.

Specify Compute Details

Specify the compute (vCenter) details for this workload domain.

Procedure

- 1 On the Compute page of the wizard, type the vCenter IP address and DNS name.
- 2 Type the vCenter subnet mask and default gateway.
- 3 Type and re-type the vCenter root password.
- 4 Click **Next**.

Select NSX Platform and Provide NSX Details

Select the NSX platform for this workload domain and provide corresponding details. The default platform is NSX-V.

Procedure

- 1 On the Networking page of the wizard, select the NSX platform.
- 2 For NSX-T, enter the VLAN ID for the Geneve overlay network.
For NSX for vSphere, enter the VLAN ID for VXLAN Networking.

Note This is the VXLAN VLAN of the management domain. A DHCP server must be configured to lease IPs in the specified VLAN. When NSX creates VXLAN VTEPs, they are assigned IP addresses from the DHCP server.

3 Provide NSX Manager details per the guidelines below.

- For NSX for vSphere in Cloud Foundation, provide the following NSX Manager details:
 - IP address
 - Name
 - Subnet mask
 - Default gateway
 - Admin password
 - Enable password (only for NSX-V)
- For NSX-T in Cloud Foundation, provide the following NSX Manager details:
 - Cluster Virtual IP address and FQDN
 - Three IP addresses and the corresponding FQDN
 - Subnet mask
 - Default gateway
 - Admin password

4 For NSX for vSphere in Cloud Foundation , provide the following NSX Controller details:

- IP addresses for the three controllers
- Subnet mask
- Default gateway
- Password

NSX-T deployment in Cloud Foundation does not include Controllers.

5 Click **Next**.

Select the vSAN Parameters

At the Storage step of the creation wizard, specify the availability you want provisioned for the VI workload domain. This page appears only if you are using vSAN storage for this workload domain.

Based on your selections, SDDC Manager will determine:

- The minimum number of hosts that it needs to fulfill those selections
- Which specific hosts in your environment are available and appropriate to fulfill those selections

- The virtual infrastructure features and their specific configurations that are needed to fulfill those selections

Note You can modify the vSAN configuration in vSphere without negatively affecting the Cloud Foundation configuration.

Procedure

- 1 Specify the level of availability you want configured for this virtual environment.

The availability level determines the level of redundancy that is set for the assigned resources. For more information, see *Managing Fault Domains in Virtual SAN Clusters in Administering VMware Virtual SAN*.

Option	Description
0	<p>With this choice, the following vSAN parameters are used:</p> <ul style="list-style-type: none"> ■ Number of failures to tolerate: zero (0). <p>Because vSAN requires a minimum of three hosts by default, three hosts are assigned to the virtual infrastructure.</p>
1	<p>With this choice, the following vSAN parameters are used:</p> <ul style="list-style-type: none"> ■ Number of failures to tolerate: one (1). <p>Because vSAN requires a minimum of three hosts by default, three hosts are assigned to the virtual infrastructure. This is the default value.</p>
2	<p>With this choice, the following vSAN parameters are used:</p> <ul style="list-style-type: none"> ■ Number of failures to tolerate: two (2). <p>Because vSAN requires a minimum of five hosts by default for this setting, five hosts are assigned to the virtual infrastructure.</p>

- 2 Click **Next**.

Specify the VMFS on FC Datastore

If you are using VMFS on FC storage for the workload domain, you must specify the VMFS on FC datastore name.

Procedure

- 1 On the Storage page, enter the name of the VMFS on FC datastore.
- 2 Click **Next**.

Select Hosts

The Host Selection page displays available hosts along with hosts details. Hosts that are powered off, cannot be accessed via SSH, or have not been properly commissioned are not displayed.

- Select only healthy hosts.

To check a host's health, SSH in to the SDDC Manager VM using the **vcf** administrative user account and type the following command:

```
sudo /opt/vmware/sddc-support/sos --health-check
```

When prompted, enter the **vcf** user password. For more information, see [Chapter 19 Supportability and Serviceability \(SoS\) Utility](#)

- For optimum performance, you must select hosts that are identical in terms of memory, CPU types, and disks.

If you select unbalanced hosts, the UI displays a warning message, but you can proceed with the workload domain creation.

- You cannot select hosts that are in a dirty state. A host is in a dirty state when it has been removed from a cluster in a workload domain.

To clean a dirty host, see [Clean up a Decommissioned Host Using the Direct Console User Interface](#).

- All selected hosts must be associated with the same network pool.

Note The Cloud Foundation API supports adding hosts from different network pools to NSX-T workload domains, as long as those network pools have the same VLAN ID and MTU settings.

Procedure

- 1 Select the hosts for creating the VI workload domain.

For a vSAN VI workload domain with 0 or 1 availability, a minimum of three hosts is required. For a VI workload domain with 2 availability, a minimum of five hosts is required. When you select hosts with sufficient storage to form a VI cluster, the **Next** button is enabled.

Starting with Cloud Foundation 3.9.1, you can add add hosts with multiple active pNICs to the workload domain by updating the API spec. For information on multi-pNIC support, see [Separating Traffic by Using Multiple vDSes](#).

The total resources based on the selected hosts are displayed.

- 2 Click **Next**.

Specify NFS Storage Details

If you are using NFS storage for this workload domain, you must provide the NFS share folder and IP address of the NFS server.

Procedure

- 1 On the NFS Storage page, enter a name for the NFS datastore name.
- 2 Enter the path to the NFS share.

- 3 Enter the IP address of the NFS server.

Note When creating additional datastores for an NFS share and server, use the same datastore name. If you use a different datastore name, vCenter overwrites the datastore name provided earlier.

- 4 Click **Next**.

Select Licenses

The Licenses page displays the available licenses for vCenter, vSphere, vSAN, and NSX based on the information you provided.

Prerequisites

You must have specified valid license keys for the following products:

- vSAN (if using vSAN as the storage option)
NFS does not require a license
- NSX for vSphere (for NSX-V VI workload domains) or NSX-T (for NSX-T VI workload domains)
- vSphere
Since vSphere and vSAN licenses are per CPU, ensure that you have sufficient licenses for the ESXi hosts to be used for the workload domain.

For information on adding license keys, see [Add License Keys for the Software in Your Cloud Foundation System](#).

Procedure

- 1 Depending on the storage option and NSX platform being used, select the appropriate licenses to apply to the VI workload domain.
- 2 Click **Next**.

View Object Names

The Object Names page displays the vSphere objects that will be generated for the VI workload domain. Object names are based on the VI workload domain name.

Procedure

- 1 Review the syntax that will be used for the vSphere objects generated for this domain.
- 2 Click **Next**.

Review Details and Start the Creation Workflow

At the Review step of the wizard, review the information about the workload domain and start the creation workflow. You can also print the information or download a printable version to print later. It can take up to two hours for the domain to be created.

The Review page displays information about the resources and their configurations that will be deployed when the workflow creates and deploys the virtual infrastructure for this workload domain.

The hosts that will be added to the workload domain are listed along with information such as the network pool they belong to, memory, CPU, and so on.

Procedure

- 1 Scroll down the page to review the information.
- 2 Click **Finish** to begin the creation process.

The Workload Domains page appears and a notification is displayed letting you know that VI workload domain is being added. Click **View Task Status** to view the domain creation tasks and sub tasks.

If a task fails, you can fix the issue and re-run the task. If the workload domain creation fails, contact VMware Support.

Results

When the VI workload domain is created, it is added to the workload domains table.

The default end-state for an NSX-T workload domain includes a network fabric with the following configuration

- two transport zones (one for VLAN and one for overlay) for each host
- NIOC and uplink profiles
- four logical switches (one each for management, vSAN, vMotion, and overlay)
- Some empty vDS and portgroups

Do not edit or delete these.

What to do next

Enable vRealize Log Insight logging for the workload domain (if not done already).

You must deploy and configure NSX-T Edge nodes manually. See [Deploy and Configure NSX Edges in Cloud Foundation](#) .

Deploy and Configure NSX Edges in Cloud Foundation

For an NSX-T VI workload domain, NSX Edges are required to enable overlay VI and public networks for north-south traffic. A minimum of two NSX Edge nodes are required for high availability and redundancy.

This document describes how to deploy NSX Edges on a VI workload domain. If you have latency intensive applications in your environment, you can deploy NSX Edges on bare metal servers. See [Deployment of VMware NSX-T Edge Nodes on Bare-Metal Hardware for VMware Cloud Foundation 3.10](#).

Sample Values for Deploying NSX-T Edges

This section lists sample values used in the NSX-T Edge procedures. These sample values have been referred through the procedures so that you can understand the network mapping for NSX-T workload domains and NSX-T Edges.

Sample Workload Domain Values

Setting	Value
Site	sfo01
NSX-T workload name	w
domain name	rainpole.local

Sample FQDN and IP Addresses for Compute vCenter Server in the Management Domain

Setting	Value	IP Address
FQDN	sfo01w01vc01.sfo01.rainpole.local	172.16.11.64

Sample FQDN and IP Addresses for NSX-T Manager in the Management Domain

FQDN	IP Address
sfo01w01nsx01a.sfo01.rainpole.local	172.16.11.82
sfo01w01nsx01b.sfo01.rainpole.local	172.16.11.83
sfo01w01nsx01c.sfo01.rainpole.local	172.16.11.84
sfo01w01nsx01.sfo01.rainpole.local	172.16.11.81

Sample VLAN IDs and IP Subnets for Availability Zone 1 for NSX-T Workload Domain

VLAN Function	VLAN ID	Subnet	Gateway
Management	1641	172.16.41.0/24	172.16.41.253
vSphere vMotion	1642	172.16.42.0/24	172.16.42.253
vSAN	1643	172.16.43.0/24	172.16.43.253
Host overlay	1644	172.16.44.0/24	172.16.44.253
UplinkO1	1647	172.16.47.0/24	172.16.47.253
UplinkO2	1648	172.16.48.0/24	172.16.48.253
Edge overlay	1649	172.16.49.0/24	172.16.49.253

Create Profiles

You must create an overlay profile, two uplink profiles (one for Edge uplink), and an Edge cluster profile.

Create Uplink Profiles

Create two uplink profiles with the failover order teaming policy with one active uplink and no standby uplinks for edge virtual machine uplink traffic.

Procedure

- 1 On the NSX-T Manager UI, navigate to **System > Fabric > Profiles**
- 2 On the Uplink Profiles tab, click **Add**.
- 3 Type the following values for the first uplink profile.

Name	Teaming Policy	Active Uplink	Transport VLAN	MTU
sfo01-w-uplink01-profile	Failover Order	Fp-eth1	1647	9000

- 4 Click **Add**.
- 5 Repeat steps 3 and 4 for the second uplink.

Name	Teaming Policy	Active Uplink	Transport VLAN	MTU
sfo01-w-uplink02-profile	Failover Order	Fp-eth2	1648	9000

Create an NSX-T Edge Cluster Profile

You must create a cluster of NSX-T Edge nodes for connectivity to the external networks and for routing availability. Before creating the Edge node cluster, define a common configuration for the nodes by creating a cluster profile.

Though an Edge cluster profile is created by default during the NSX-T workload domain creation, you must create a new cluster profile so that you can modify it if required.

Procedure

- 1 On the NSX-T Manager UI, navigate to **System > Fabric > Profiles**.
- 2 On the Edge Cluster Profiles tab, click **Add**.
- 3 Type the following values. These example values are built upon the [Sample Values for Deploying NSX-T Edges](#).

Setting	Example Value
Name	sfo01-w-edge-cluster01-profile
BFD Probe	1000

Setting	Example Value
BFD Allowed Hops	255
BFD Declare Dead Multiple	3

- 4 Click **Add**.

Create Transport Zones

Two transport zones (one each for overlay and VLAN traffic) are created during the NSX-T workload domain creation process. You must create two additional transport zones - one for each Edge uplink.

Procedure

- 1 On the main navigation bar of the NSX-T Manager UI, click **System**.
- 2 Navigate to **Fabric > Transport Zones** and click **Add**.
- 3 Enter the following values.

Setting	Example Value
Name	sfo01-w-uplink01
N-VDS Name	sfo01-w-uplink01
N-VDS Mode	Standard
Traffic Type	VLAN

- 4 Click **Add**.
- 5 Repeat step 4 to create the second transport zones with the following values.

Setting	Example Value
Name	sfo01-w-uplink02
N-VDS Name	sfo01-w-uplink02
N-VDS Mode	Standard
Traffic Type	VLAN

- 6 Click **Add** and **Save**.

Create NSX-T Segments for System, Uplink, and Overlay Traffic

Three segments (one each for management, vMotion, and vSAN networks) are created in the VLAN transport zone by Cloud Foundation during the NSX-T workload domain creation process. You must create additional segments manually for both Edge uplinks before deploying the Edge VMs to connect nodes that send VLAN and overlay traffic. If you plan to deploy the Enterprise PKS solution, you must also create a segment for Enterprise PKS management and service networks.

Prerequisites

The table below lists sample names for segments you need to create as well as sample VLAN IDs.

Segment Name	Transport Zone	VLAN
sfo01-w-nvds01-uplink01	sfo01-esxi-vlan	0-4094
sfo01-w-nvds01-uplink02	sfo01-esxi-vlan	0-4094
sfo01-w-uplink01	sfo01-w-uplink01	1647
sfo01-w-uplink02	sfo01-w-uplink02	1648

Procedure

- 1 Log in to the NSX-T Manager UI with the `admin` user name and `comnsx_admin_password`.
- 2 Navigate to **Networking > Segments**.
- 3 On the Segments tab, click **Add Segments**.
- 4 Create a segment with the following sample values.

Setting	Sample Value
Segment Name	sfo01-w-uplink01
Connected Gateway & Type	None
Transport Zone	sfo01-w-uplink01
VLAN	1647

- 5 Repeat steps 3 and 4 for each segment to be created.
- 6 Create the following logical switches if you are deploying the Enterprise PKS solution. These must be created through the Advanced Networking and Security tab.

Segment Name	Transport Zone	VLAN	Type
pks-management	overlay-tz-ID	overlay	N/A
pks-service	overlay-tz-ID	overlay	N/A

Deploy NSX-T Edge Appliances

Deploy two NSX-T Edge appliances to provide workload domains with routing services and connectivity to external networks.

Table 8-4. Example Values for NSX-T Nodes

Setting	Example Values for Node 1	Example Values for Node 2
Name	sfo01wesg01	sfo01wesg02
	ls-Cluster ID-management	ls-Cluster ID-management
Primary IP Address	172.16.41.21	172.16.41.22

Prerequisites

Download the NSX Edge appliance OVA file (called NSX Edge for VMware ESXi) on the VMware download portal. Either copy the download URL or download the OVA file onto your computer.

Procedure

- 1 In a browser window, navigate to the vCenter Server UI by typing `https://vCenter of WLD/ui`.
- 2 Login with the `administrator@vsphere.local` username and admin password.
- 3 In the Hosts and Clusters inventory, expand the cluster where you want to place the Edge appliances.

Edge appliances must be placed in a workload domain cluster, not the management domain cluster.

- 4 Right-click the cluster and select **Deploy OVF Template**.
- 5 Navigate to the NSX Edge appliance OVA file and click **Next**.
- 6 Type a name for the appliance. For example, `sfo01wesg01`.
- 7 Click **Next**.
- 8 Select the cluster and click **Next**.
- 9 Select `vsan` and click **Next**.
- 10 Type the following information.

Source Network	Example Destination Network
Network 3	<code>sfo01w-nvds01-uplink02</code>
Network 2	<code>sfo01w-nvds01-uplink01</code>
Network 0	<code>ls-Cluster ID-management</code>

The overlay network segment is generated by Cloud Foundation during workload domain creation. This segment is named `net-overlay-UUID`.

- 11 Click **Next**.
- 12 On the Customize template page, expand the setting groups, and enter the following values.

Application

Setting	Example Value
System Root User Password	<code>nsx_edge_root_password</code>
CLI "admin" User Password	<code>nsx_edge_admin_password</code>
CLI "audit" User Password	<code>nsx_edge_admin_password</code>

Network Properties

Setting	Example Value
Host name	sfo01wesg01.sfo01.rainpole.local
Default IPv4 Gateway	172.16.41.253
Management Network IPv4 Address	172.16.41.21
Management Network Netmask	255.255.255.0

DNS and Services Configuration

Setting	Example Value
DNS Server List	172.16.11.5 172.16.11.4
Domain Search List	rainpole.local
NTP Server List	172.16.11.251
Enable SSH	Select
Allow root SSH login	Do not select

- 13 Click **Finish**.
- 14 After the deployment is completed, power on the NSX-T Edge appliance.
- 15 In the VMs and Templates inventory, find the sfo01wesg01 virtual machine, and select **Power > Power On**.
- 16 Repeat steps 3 - 14 to deploy the second Edge node (sfo01wesg02). Refer to the sample values at the beginning of this topic.
- 17 Create a resource pool for the Edge appliance VMs.
 - a Right-click the sfo01-w02-shared01 cluster and select **New Resource Pool**
 - b In the New Resource Pool dialog box, enter the values for the sfo01-w02rp-sddc-edge resource pool and click **OK**.

Setting	Example Value
Name	sfo01-w02rp-sddc-edge
CPU-Shares	High
CPU-Reservation	0
CPU-Reservation Type	Expandable selected
CPU-Limit	Unlimited
Memory-Shares	Normal
Memory-Reservation	32 GB
Memory-Reservation Type	Expandable selected
Memory	Unlimited

- 18 Create a folder for the NSX-T Edge VMs for inbound and outbound traffic in the workload domain.
 - a From the **Home** menu, select **VMs and Templates**.
 - b In the inventory tree, expand `sfo01w02vc01.sfo01.rainpole.local`.
 - c Right-click the `sfo01-w02dc` data center and select **New Folder > New VM and Template Folder > .**
 - d In the New Folder dialog box, enter a name for the folder (such as `sfo01-w02fd-nsx`) and click **OK**.
 - e In the New Resource Pool dialog box, enter the values for the `sfo01-w02rp-sddc-edge` resource pool and click **OK**.

Join the NSX-T Edge Nodes to the Management Plane

Join the NSX-T Edge nodes to the NSX-T management cluster.

Setting	Sample Value for sfo01wesg01	Sample Value for sfo01wesg02
Port Groups	ls- Cluster ID -management	ls- Cluster ID -management
Primary IP Address	172.16.41.21	172.16.41.22

Procedure

- 1 Log in to the first NSX-T Manager node using Secure Shell (SSH) client.

Setting	Value
FQDN	NSX-T Manager IP address or FQDN
User name	admin
Password	<i>nsx_admin_password</i>

- 2 Run the following command to retrieve the thumbprint ID for the NSX-T Manager cluster.
`get certificate cluster thumbprint`
- 3 Copy the output. You will use this output in step 5.
- 4 Log in to the first NSX-T Edge (`sfo01wesg01`) node using Secure Shell (SSH) client.

Setting	Value
FQDN	<code>sfo01wesg01.??</code>
User name	admin
Password	<i>edge_admin_password</i>

- 5 Join the NSX-T Edge node to the management plane by running the following command.

```
join management-plane sfo01wnsx01.sfo01.rainpole.local thumbprint thumbprintid username admin
```

where *thumbprintid* is the output you copied in step 3.

- 6 Enter the password for the admin account.
- 7 Repeat steps 3 - 5 on the second Edge node (sfo01wesg02).

Configure Edge Nodes as Transport Nodes

NSX-T Edges must be configured as transport nodes so that they can join the other transport nodes for Overlay traffic. Creating them as transport nodes sets the TEPs and N-VDSes on the Edges.

Procedure

- 1 In a web browser, navigate to the NSX-T UI and log in with the following credentials.

Setting	Value
User name	admin
Password	<i>compnsx_admin_password</i>

- 2 On the navigation bar, click **System**.
- 3 In the navigation plane, select **Fabric > Nodes > Edge Transport Nodes**.
- 4 Under Edit Transport Node - sfo01wesg01, click the General tab.
- 5 Under Transport Zones, move the following transport zones to the Selected list.
 - sfo01-w-uplink01(VLAN)
 - sfo01-w-uplink02(VLAN)
 - overlay-tz-<ID>(Overlay)
- 6 Click **Add**.
- 7 Under Edit Transport Node - sfo01wesg01, click the N-VDS tab.
- 8 Under New Node Switch, enter the following switch configuration.

Setting	Value for sfo01wesg01
Edge Switch Name	sfo01wuplink01(VLAN)
Uplink Profile	sfo01woverlay-profile
IP Assignment	Use static IP list
Static IP List	172.16.49.21
Gateway	172.16.49.253
Subnet Mask	255.255.255.0
Virtual NICs	fp-eth0, sfo01-w-uplink01-profile

- 9 Click **Add N-VDS** and provide the required information. See the table below for sample values.

Setting	Sample Values for Edge Node1
Edge Switch Name	Edge-uplink1-nvds
Uplink Profile	sfo01-w-uplink01-profile
IP Assignment	Not applicable
Virtual NICs	fp-eth1, sfo01-w-uplink01-profile

- 10 Repeat steps 4 - 9 for the second Edge node (sfo01wesg02 in our example).

Setting	Sample Values for Edge Node1
Edge Switch Name	Edge-uplink2-nvds
Uplink Profile	sfo01-w-uplink02-profile
IP Assignment	Not applicable
Virtual NICs	fp-eth2, sfo01-w-uplink01-profile

Results

The Edge nodes have the following values..

Setting	Value for sfo01wesg01	Value for sfo01wesg02
Edge	sfo01wesg01	sfo01wesg02
Management IP	172.16.41.21	172.16.41.22
Node Status	Up	Up
Transport Zones	<ul style="list-style-type: none"> ■ sfo01-w-uplink01 ■ sfo01-w-uplink02 ■ overlay-tz-<ID> 	<ul style="list-style-type: none"> ■ sfo01-w-uplink01 ■ sfo01-w-uplink02 ■ overlay-tz-<ID>
N-VDS	Depends on the number of N-VDses added in step 9.	3

Create an NSX-T Edge Cluster

Add both NSX-T Edge nodes to a cluster to increase the availability of networking services.

Procedure

- 1 In a web browser, navigate to the NSX-T UI and log in with the following credentials.

Setting	Value
User name	admin
Password	compnsx_admin_password

- 2 On the navigation bar, click **System**.

- 3 In the navigation plane, select **Fabric > Nodes**.
- 4 Click **Add**.
- 5 Configure the following settings.

Setting	Sample Values
Name	sfo01w-edge-cluster01
Edge Cluster Profile	sfo01w-edge-cluster01-profile

- 6 From the Member Type drop-down menu, select **Edge Node**.
- 7 Move the sfo01wesg01.sfo01.rainpole.local and sfo01wesg02.sfo01.rainpole.local nodes to the the Selected list.
- 8 Click **OK** and then click **Add**.

Configure Dynamic Routing

Configure dynamic routing to enable communication between tenant workloads by using application virtual networks in NSX-T and to connect tenant workloads to the external network.

Routing needs to be configured in north-south and east-west direction.

- North-South traffic leaving or entering the workload domain, for example, a virtual machine on an overlay network communicating with an end-user device on the corporate network.
- East-West traffic remains in the workload domain, for example, two virtual machines on the same or different segments communicating with each other.

Procedure

1 [Add a Tier-0 Gateway](#)

Add and configure a tier-0 gateway in the NSX-T Edge cluster to provide a gateway service between the logical and physical network.

2 [Add Tier-1 Gateway](#)

Add and configure a tier-1 gateway to re-distribute routes to the tier-0 gateway and provide routing between tenant workloads. This is required only when you enable a centralized service, like LB or NAT. Tier-1 gateways have downlink ports to connect to NSX-T segments and uplink ports to connect to NSX-T tier-0 gateways.

3 [Verify BGP Peering and Route Redistribution](#)

The Tier-0 gateway must establish a connection to each of the upstream Layer 3 devices before BGP updates can be exchanged. Verify that the NSX-T Edge nodes are successfully peering and that BGP routing is established.

Add a Tier-0 Gateway

Add and configure a tier-0 gateway in the NSX-T Edge cluster to provide a gateway service between the logical and physical network.

Procedure

- 1 In a browser window, log in with admin privileges to the NSX-T Manager at <https://nsx-manager-ip-address>.
- 2 On the main navigation bar, click **Advanced Networking & Security**.
- 3 Select **Networking > Routers** and click **Add Tier-0 Router**.
- 4 Type the following values and click Save.

Setting	Value
Name	sfo01-w-tier-0-01
High Availability Mode	Active-Standby
Edge cluster	sfo01-w-edge-cluster01

- 5 Select sfo01-w-tier-0-01 and click **Edit**.
- 6 Select **Routing > Route redistribution**
- 7 Configure route redistribution.
 - a Click **Edit**.
 - b Select **Enable**.
 - c Click **Add** to add new redistribution criteria.
 - d Select all sources and click **Apply**.
- 8 Select **Configuration > Route ports**.
- 9 Click **Add** and enter the settings of the uplink interfaces.

Name	Type	IP Address / Mask	Connected to (Segment)	Edge Node	MTU
sfo01wesg01-Uplink01	External	172.16.47.2/24	sfo01-w-uplink01	sfo01wesg01	9000
sfo01wesg01-Uplink02	External	172.16.48.3/24	sfo01-w-uplink02	sfo01wesg02	9000
sfo01wesg01-Uplink03	External	172.16.47.4/24	sfo01-w-uplink03	sfo01wesg03	9000
sfo01wesg01-Uplink04	External	172.16.48.5/24	sfo01-w-uplink04	sfo01wesg04	9000

10 Configure BGP.

- a Expand **Routing > BGP** and click Edit BGP Configuration.
- b Enter the following settings.

Setting	Value
Local AS	65000
BGP	On
Graceful Restart	helper-mode
Inter SR iBGP	On
ECMP	Off
Multipath Relax	Off

- c In BGP Networks, click **Add**.
- d Add BGP Neighbor and specify the following settings for the first layer 3 device.

IP Address	BFD	Remote AS	Hold Down Time	Keep Alive Time	Password
172.16.47.1	Disabled	65001	12	4	<i>bgp_password</i>
172.16.48.1	Disabled	65001	12	4	<i>bgp_password</i>

11 Click **Close Editing**.**12** Generate a BGP summary for the tier-0 gateway.

- a Select **Routers** and select sfo01-w-tier-0-01.
- b Click **Actions > Generate BGP Summary > .**
- c Verify that the connection status of each transport node is **Established**.

Add Tier-1 Gateway

Add and configure a tier-1 gateway to re-distribute routes to the tier-0 gateway and provide routing between tenant workloads. This is required only when you enable a centralized service, like LB or NAT. Tier-1 gateways have downlink ports to connect to NSX-T segments and uplink ports to connect to NSX-T tier-0 gateways.

Procedure

- 1** In a browser window, log in with admin privileges to the NSX-T Manager at <https://nsx-manager-ip-address>.
- 2** On the main navigation bar, click **Advanced Networking & Security**.
- 3** Select **Networking > Routers** and click **Add Tier-1 Router**.

4 Type the following values and click Save.

Setting	Value
Name	sfo01-w02-tier-1-01
Linked Tier-0 Gateway	sfo01-w-tier-0-01
Edge Cluster	sfo01-w-edge-cluster-01

5 Next to Edges, click **Set**.

6 Click **Add Edge**.

7 Click sfo01wesg01 and sfo01wesg02 edge nodes and click **Apply**.

8 Select the newly created tier-1 router (sfo01-w02-tier-1-01).

9 Verify the connection between tier-1 and tier-0 gateways.

- a Click **Routers** and then select the sfo01-w02-tier-1-01 router.
- b Click **Configuration > Router Ports** and verify that the LinkedPort has the following settings.

Setting	Expected Value
Logical Router	LinkedPort_sfo01-w-tier-0-01
Type	Linked Port
IP Address/mask	x.x.x.x/31
Connected to	sfo01-w-tier-0-01
Transport node	sfo01wesg01 sfo01wesg02

10 Select the sfo01-w-tier-0-01 gateway.

11 Click **Configuration > Router Ports** and verify that the LinkedPort has the following settings.

Setting	Expected Value
Logical Router	LinkedPort_sfo01-w-tier-0-01
Type	Linked Port
IP Address/mask	x.x.x.x/31
Connected to	sfo01-w-tier-0-01
Transport node	sfo01wesg01 sfo01wesg02

Results

Verify BGP Peering and Route Redistribution

The Tier-0 gateway must establish a connection to each of the upstream Layer 3 devices before BGP updates can be exchanged. Verify that the NSX-T Edge nodes are successfully peering and that BGP routing is established.

Procedure

- 1 Using SSH, log in to the first Edge node (sfo01wesg01 with the following credentials:

User name: admin

Password: *nsx_edge_admin_password*

- 2 Get information about the Tier-0 and Tier-1 service routers and distributed router.

```
get logical-router
```

For example, the output of the command might contain the following configuration:

UUID	VRF	LR-ID	Name	Type	Ports
sample_uuid	0	0		TUNNEL	3
sample_uuid	1	5	SR-tier0-01	SERVICE_ROUTER_TIER0	6
sample_uuid	2	2	DR-tier1-01	DISTRIBUTED_ROUTER_TIER1	5
sample_uuid	3	3	DR-tier0-01	DISTRIBUTED_ROUTER_TIER0	4
sample_uuid	4	11	SR-tier1-01	SERVICE_ROUTER_TIER1	5

- 3 Using the VRF value for SERVICE_ROUTER_TIER0, connect to the service router for Tier-0.

```
vrf 1
```

- 4 The prompt changes to `hostname(tier0_sr)>`. All commands are associated with this object.
- 5 Verify the BGP connections to the neighbors of the service router for Tier-0.

```
get bgp neighbor
```

The BGP State for each neighbor appears as Established, up for hh:mm:ss.

- 6 Verify that you are receiving routes by using BGP and that multiple routes to BGP-learned networks exist.

```
get route
```

- 7 Repeat this procedure on the second Edge node.

View Workload Domain Details

The Workload Domains page displays high level information about the workload domains in the Cloud Foundation system. CPU, memory, and storage utilized by the workload domain is also displayed here.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
- 2 In the workload domains table, click the name of the workload domain.

The domain details page displays CPU, memory, and storage allocated to the domain. The tabs on the page display additional information as described in the table below.

Tab	Information Displayed
Summary	Clusters in the workload domain and availability level for each cluster.
Services	SDDC software stack components deployed for the workload domain's virtual environment and their IP addresses. Click a component name to navigate to that aspect of the virtual environment. For example, click vCenter to reach the vSphere Web Client for that workload domain. All the capabilities of a VMware SDDC are available to you in the VI workload domain's environment, such as creating, provisioning, and deploying virtual machines, configuring the software-defined networking features, and so on.
Updates/Patches	Available updates for the workload domain.
Update History	Updates applied to this workload domain.
Hosts	Names, IP addresses, status, associated clusters, and capacity utilization of the hosts in the workload domain and the network pool they are associated with.
Clusters	Names of the clusters, number of hosts in the clusters, and their capacity utilization.
Security	Default certificates for the Cloud Foundation components. For more information, see Chapter 4 Certificate Management .

What to do next

You can add a cluster to the workload domain from this page.

Delete a VI Workload Domain

When you delete a workload domain, the clusters within that workload domain are deleted and the hosts are returned to the free pool.

Monitoring through Log Insight and vRealize Operations is removed and the components associated with the workload domain to be deleted contained within the management domain are removed. This includes the vCenter Server instance and NSX Manager.

The network pools used by the workload domain are not deleted as part of the workload domain deletion process and must be deleted separately.

Caution Deleting a workload domain is an irreversible operation. All clusters and VMs within the workload domain are deleted and the underlying datastores are destroyed.

It can take up to 20 minutes for a workload domain to be deleted. During this process, you cannot perform any operations on workload domains.

Prerequisites

- Back up the data on the workload domain. The datastores on the workload domain are destroyed when the workload domain is deleted.
- Migrate the VMs that you want to keep to another workload domain.
- For an NSX-T workload domain, delete any workload VMs created outside Cloud Foundation before deleting the workload domain.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 Hover your mouse in the workload domain row that where you want to delete.

When you select the workload domain, three vertical dots appear next to the name.

- 3 Click the dots and choose **Delete Domain**.

A confirmation window appears with details about the impact of deleting the workload domain, including how many hosts will be returned to the free pool.

- 4 Click **Delete Domain** to proceed.

The details page for the workload domain appears with a message indicating that the workload domain is being deleted. When the removal process is complete, the workload domain is removed from the domains table.

View Cluster Details

The cluster page displays high level information about the cluster as well as the hosts that form that cluster. CPU, memory, and storage utilization for this cluster is also displayed here.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domain**.
- 2 In the workload domains table, click the name of a workload domain.
- 3 Click the **Clusters** tab.

- 4 In the clusters table, click the name of a cluster.

The cluster detail page appears. The tabs on the page display additional information as described in the table below.

Tab	Information Displayed
Summary	Storage parameter configured on the cluster and organization name.
Hosts	Details about each host in the cluster. You can click a name in the FQDN column to access the host detail page.

What to do next

You can add or remove a host, or access the vSphere Client from this page.

Shrink a Workload Domain

You can reduce the management domain or a VI workload domain by removing a host from a vSphere cluster in the workload domain or by deleting a vSphere cluster.

Remove a Host from a Cluster in a Workload Domain

You can remove a host from a cluster in a workload domain through the Workload Domains page in the SDDC Manager Dashboard.

When a host is removed, the vSAN members are reduced. Ensure that you have enough hosts remaining to facilitate the configured vSAN availability. Failure to do so might result in the datastore being marked as read-only or in data loss.

Prerequisites

If you are removing a host from an NSX-T workload cluster, delete any workload VMs created outside Cloud Foundation or move them to another host.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
The Workload Domains page displays information for all workload domains.
- 2 In the workload domains table, click the name of the workload domain that you want to modify.
The detail page for the selected workload domain appears.
- 3 Click the **Clusters** tab.
- 4 Click the name of the cluster from which you want to remove a host.
- 5 Click the **Hosts** tab.

6 Select the host to remove and click **Remove Selected Hosts**.

An alert appears, asking you to confirm or cancel the action. If the removal results in the number of hosts in the cluster being less than the minimum number of required hosts, you must click **Force Remove** to remove the host.

7 Click **Remove** to confirm the action.

The details page for the cluster appears with a message indicating that the host is being removed. When the removal process is complete, the host is removed from the hosts table.

Results

The host is removed from the workload domain and added to the free pool.

What to do next

Clean up the host so that you can use it again. See [Clean up a Decommissioned Host Using the Direct Console User Interface](#).

Delete a Cluster from a Workload Domain

You can delete a cluster from a workload domain. Datastores on the ESXi hosts in the deleted cluster are destroyed.

You cannot delete the last cluster in a workload domain. Instead, delete the workload domain. See [Delete a VI Workload Domain](#).

Prerequisites

- If you are deleting an NSX-T workload cluster, delete any workload VMs created outside Cloud Foundation before deleting the cluster.

Migrate or backup the VMs and data on the data store associated with the cluster to another location.

Procedure

1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

2 Click the name of the workload domain that contains the cluster you want to delete.

3 Click the **Clusters** tab to view the clusters in the workload domain.

4 Hover your mouse in the cluster row you want to delete.

5 Click the three dots next to the cluster name and click **Delete Cluster**.

6 Click **Delete Cluster** to confirm that you want to delete the cluster.

The details page for the workload domain appears with a message indicating that the cluster is being deleted. When the removal process is complete, the cluster is removed from the clusters table.

Expand a Workload Domain

You can expand the management domain or a VI workload domain to add resources to support additional workloads or availability.

To expand a domain, you can:

- Add a host from the Cloud Foundation inventory to a vSphere cluster.
By adding an individual host to an existing workload domain, you can expand the amount of resources contained within an existing vSphere cluster.
- Add a new vSphere cluster to a workload domain.
As workload domains support multiple vSphere clusters, you can add an additional cluster to an existing workload domain to provide for increased capacity and VM failover isolation.

Add a Host to a Cluster in a Workload Domain

Adding an individual host to a workload domain adds the resources of that host to the workload domain. You can add multiple hosts at a time to a workload domain.

Prerequisites

- There must be a host available in the Cloud Foundation inventory. For information on adding a host to Cloud Foundation, see [Commission Hosts](#).
- Ensure that the host you want to add is in an active state.
- You must have a valid vSphere license specified in the Licensing tab of the SDDC Manager Dashboard with adequate sockets available for the host to be added. For more information, see [Add License Keys for the Software in Your Cloud Foundation System](#).
- Verify that the host to be added to the workload domain matches the configuration of the hosts in the cluster to which you want to add the domain. This allows the cluster configuration to remain balanced. If the host to be added does not match the pre-existing hosts in the cluster, the cluster will be unbalanced and a warning will be displayed. The warning will not prevent the expansion and can be dismissed if needed.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
The Workload Domains page displays information for all workload domains.
- 2 In the workload domains table, click the name of the workload domain that you want to expand.
The detail page for the selected workload domain appears.
- 3 Click the **Clusters** tab.
- 4 Click the name of the cluster where you want to add a host.

5 Click **Actions > Add Host**.

The Add Hosts wizard appears.

6 Select the host you want to add to the cluster.

The host you select must be associated with the same network pool as the other hosts in the cluster.

Note The Cloud Foundation API supports adding hosts from different network pools to NSX-T workload domain clusters, as long as those network pools have the same VLAN ID and MTU settings.

For optimum performance, you should select hosts that are identical in terms of memory, CPU types, and disks to the other hosts in the cluster. If you select unbalanced hosts, the UI displays a warning message, but you can proceed with the workload domain creation.

Starting with Cloud Foundation 3.9.1, you can add a host with multiple active pNICs to a multi pNIC-aware cluster using APIs. For information on multi-pNIC support, see [Separating Traffic by Using Multiple vDSes](#).

7 Click **Next**.**8** Select the vSphere license you want to apply to the host.**9** Click **Next**.**10** Review the host and license details, and click **Finish**.

The details page for the cluster appears with a message indicating that the host is being added. Wait until the action is complete before performing additional workload domain tasks.

Add a Cluster to a Workload Domain

You can add a cluster to a workload domain through the Workload Domains page in the SDDC Manager Dashboard.

Prerequisites

- There must be at least three hosts available in the Cloud Foundation inventory. For information on adding a host to Cloud Foundation, see [Commission Hosts](#).
- Ensure that the hosts you want to add to the cluster are in an active state.
- You must have a valid vSphere and vSAN (if using vSAN storage) license specified in the Licensing tab of the SDDC Manager Dashboard with adequate sockets available for the host to be added. For more information, see [Add License Keys for the Software in Your Cloud Foundation System](#).
- A DHCP server must be configured on the VXLAN VLAN of the management domain. When NSX creates VXLAN VTEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.

- From the UI, you can only select hosts on which pNICs 0 and 1 are active. For information on multi-pNIC support, see [Separating Traffic by Using Multiple vDSes](#).

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 Use one of the following methods to get started.

- ◆ From the high level workload domain page:

- a Hover your mouse in the workload domain row that where you want to add a cluster.

A set of three dots appear on the left of the workload domain name.

- b Click these dots and then click **Add Cluster**.

The Add Cluster wizard appears.

- ◆ From the workload domain details page:

- a Click the name of the workload domain to go to the details page for that workload domain.

- b Click **Actions > Add Cluster**.

The Add Cluster wizard appears.

- 3 Select the storage type for the cluster and click **Begin**.

- 4 Enter a name for the cluster and click **Next**.

- 5 On the Networking page, enter the VXLAN VLAN of the management domain and click **Next**.

This is the VXLAN VLAN of the management domain. A DHCP server must be configured to lease IPs in the specified VLAN. When NSX creates VXLAN VTEPs, they are assigned IP addresses from the DHCP server.

- 6 If you selected vSAN storage for the cluster, the vSAN parameters page appears. Specify the level of availability you want configured for this cluster. The specified Failures To Tolerate (FTT) value determines the number of hosts required the cluster.

- 7 If you selected VMFS on FC storage for the cluster, enter the VMFS on FC datastore name.

- 8 Click **Next**.

- 9 On the Object Names page, review the syntax that will be used for the vSphere objects generated for this cluster and click **Next**.

- 10 On the Host Selection page, select hosts for the cluster.

All hosts in a cluster must be associated with the same network pool.

Note The Cloud Foundation API supports adding hosts from different network pools to NSX-T workload domain clusters, as long as those network pools have the same VLAN ID and MTU settings.

Starting with Cloud Foundation 3.9.1, you can add add hosts with multiple active pNICs to the cluster using APIs. For information on multi-pNIC support, see [Separating Traffic by Using Multiple vDSes](#).

When you have selected the minimum number of hosts required for this cluster, the Next button is enabled.

- 11 Click **Next**.
- 12 If you selected NFS storage for the cluster, the NFS Storage page appears. Enter the datastore name, NFS share folder, and NFS server IP address.
- 13 Click **Next**.
- 14 On the Licenses page, select the vSphere and vSAN (if using vSAN storage) license to apply to this cluster.
- 15 Click **Next**.
- 16 On the Review page, review the cluster details and click **Finish**.

The details page for the workload domain appears with the following message: Adding a new cluster is in progress. When this process completes, the cluster appears in the Clusters tab in the details page for the workload domain.

Working with VMware Enterprise PKS

9

VMware Cloud Foundation enables automated deployment of VMware Enterprise PKS on an NSX-T workload domain.

Enterprise PKS is a container services solution that simplifies the deployment and management of Kubernetes clusters. Enterprise PKS manages container deployment from the application layer all the way to the infrastructure layer according to the requirements for production-grade software. Enterprise PKS supports high availability, autoscaling, health-checks and self-repairing of underlying virtual machines, and rolling upgrades for the Kubernetes clusters.

Enterprise PKS on Cloud Foundation is deployed on a single site. Within that site, stretched vSAN clusters are supported.

Enterprise PKS Components

The Enterprise PKS solution includes the following components:

- PKS appliance deploys VMs that form Kubernetes clusters (master and worker nodes).
- Harbor Registry is an open source enterprise cloud native registry that stores, signs, and scans container images for vulnerabilities.
- BOSH Director keeps your Kubernetes environment healthy.
- Operations Manager configures PKS appliance, Harbor Registry, and BOSH Director.

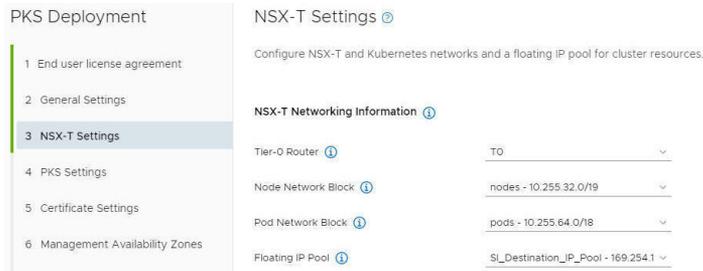
Enterprise PKS Availability Zones

Within Enterprise PKS, infrastructure is allocated into availability zones. Applications are distributed across these availability zones for high availability. There are three types of availability zones:

- Management availability zone
Enterprise PKS management components (PKS appliance, Harbor Registry, BOSH Director, and Operations Manager) are deployed in the management availability zone.
- Compute availability zones
Kubernetes cluster nodes are deployed in compute availability zones.

- **Kubernetes availability zone**

Kubernetes availability zone (also referred to as service network) is for legacy workloads and is not used in the Enterprise PKS solution deployed by Cloud Foundation.



This chapter includes the following topics:

- [Deploy an Enterprise PKS Solution](#)
- [View PKS Solution Details](#)
- [Delete a Enterprise PKS Solution](#)
- [Expand or Shrink an NSX-T Workload Domain associated with an PKS Solution](#)

Deploy an Enterprise PKS Solution

To deploy an Enterprise PKS, perform the steps below in the order in which they are documented.

Prerequisites for Deploying VMware Enterprise PKS

This sections lists the prerequisites for deploying Enterprise PKS.

- An NSX-T backed VI workload domain must be deployed.
- Shared NSX-T Edges must be deployed and the logical North-South routing must be configured. See [Deploy and Configure NSX Edges in Cloud Foundation](#) .
- Decide on the static IP addresses, host names, and subnets for the Enterprise PKS components - PKS API, BOSH Director, Operations Manager, and Harbor Registry. Verify that these IP addresses and FQDNs are allocated on the DNS server and are available for deployment. The examples in these procedures use the IP addresses defined below. Replace them with IP addresses being used in your network.

Component	FQDN	IP Address	Notes
Operations Manager	sfo01w02ops01.rainpole.local	10.0.0.10	This IP must be within the exclusion range
BOSH Director for vSphere	sfo01w02bosh01.rainpole.local	10.0.0.11	This IP should be the first one after the last IP in the exclusion range (last value in exclusion range + 1)

Component	FQDN	IP Address	Notes
Enterprise PKS	sfo01w02pkfdb01.rainpole.local	10.0.0.12	This IP should be the second one after the last IP in the exclusion range (last value in exclusion range + 2)
Harbor Container Registry	sfo01w02hrbr.rainpole.local	10.0.0.14	This IP should be the third one after the last IP in the exclusion range (last value in exclusion range + 3)

Sample subnets for Enterprise PKS components are in the table below.

Component	CIDR	Gateway	Route
PKS management workloads	10.0.0.0/24	10.0.0.1	Route to/from SDDC management network
PKS service networks	10.0.1.0/24	10.0.1.1	Route to/from SDDC management network
Nodes IP block	10.1.0.0/16	10.1.0.1	Route to/from SDDC management network and external access to docker.io.
Pods IP block	10.2.0.0/16	10.2.0.1	
Floating IP pool	10.3.0.0/24	10.3.0.1	External access

- Two IP blocks and a floating pool of IP addresses must be configured within the NSX-T domain. Enterprise PKS assigns IP addresses from this pool to VIP of the load balancers for Kubernetes Pod services.

Refer to the example IP block and IP pool tables below.

IP Block	CIDR
sfo01-w-nodes-ip-block	10.1.0.0/16
sfo01-w-pods-ip-block	10.2.0.0/16

Setting	Value
IP Ranges	10.3.0.10-10.3.0.250
CIDR	10.3.0.0/24

- Generate a certificate and key for an NSX-T Manager super user

Enterprise PKS uses a super user to create, modify, and delete objects in NSX-T. Run the script below to create a certificate and key for that super user. This certificate and key is required as input in the [Provide Certificate Details](#) page of the Deploy PKS wizard.

```
#!/bin/bash

NSX_SUPERUSER_CERT_FILE="pks-nsx-t-superuser.crt"
NSX_SUPERUSER_KEY_FILE="pks-nsx-t-superuser.key"

openssl req \
  -newkey rsa:2048 \
  -x509 \
  -nodes \
  -keyout "$NSX_SUPERUSER_KEY_FILE" \
  -new \
  -out "$NSX_SUPERUSER_CERT_FILE" \
  -subj /CN=ocp-nsx-t-superuser \
  -extensions client_server_ssl \
  -config <(
    cat /etc/ssl/openssl.cnf \
    <(printf '[client_server_ssl]\nextendedKeyUsage = clientAuth\n')
  ) \
  -sha256 \
  -days 730

cat pks-nsx-t-superuser.crt
cat pks-nsx-t-superuser.key
```

- Generate CA-Signed Certificates for Operations Manager, BOSH Director, Enterprise PKS control plane, and Harbor Registry. The certificates must include the fully qualified domain name for each component. You use these certificates for trusted communication between the Enterprise PKS components and the rest of the environment.
- Prepare the network settings and resources for availability zones.

You can achieve availability by deploying the Kubernetes cluster nodes across multiple compute availability zones. You must configure the network CIDR, gateway, reserved IP ranges, and target logical switch for the availability zones. Depending on the storage being used in your environment, there are two ways to define availability zones:

- When using NFS storage in a multi-cluster topology, availability zones are mapped to vCenter clusters.

Add three clusters to the NSX-T VI workload domain. NFS storage across clusters is required for persistent volume which is accessible by all hosts in the clusters.

- When using VMFS on FC storage in a multi-cluster topology, availability zones are mapped to vCenter clusters.

Add three clusters to the NSX-T VI workload domain. VMFS on FC storage across clusters is required for persistent volume which is accessible by all hosts in the clusters.

- When using vSAN storage in a single-cluster topology, availability zones are mapped to resource pools in the default cluster.

Create four resource pools in the default NSX-T VI workload domain cluster. One resource pool is used for the Enterprise PKS VMs and the other resource pool is used for Kubernetes work nodes.

For example, create the following resource pools in the default cluster:

- RP-SharedAZ
 - RP-Comp1AZ
 - RP-Comp2AZ
 - RP-Comp3AZ
- Download the install bundle for Enterprise PKS. See [Chapter 13 Downloading an Install Bundle](#).

Start the Deploy Wizard and Specify General Settings

On the General Settings page, select the NSX-T workload domain where you want to deploy the Enterprise PKS solution and provide the authentication settings for the management components.

Procedure

- 1 On the SDDC Manager Dashboard, select Workload Domains from the navigation panel.
- 2 In the PKS section, click Get Started.
- 3 Select the checkboxes to ensure that all required prerequisites have been met and then click **Begin**.

PKS Deployment Prerequisites

Before you deploy, complete the following required prerequisites.

- Select All
- NSX-T Workload Domain
An NSX-T backed workload domain must be prepared and available for consumption.
- DNS and IP Allocation
Prepare the IP addresses and forward/reserve DNS records for the PKS API, Pivotal Operations Manager, and the Harbor Registry (optional).
- Certificates
Generate the certificates and private keys from a trusted certificate authority that include the fully qualified domain names for each PKS management component.
- NSX-T Overlay Networking
Prepare the NSX-T Tier-0 router, node and pod IP blocks, and a floating IP pool for Kubernetes cluster resources.
- Availability Zones
Prepare the network settings and resources for the availability zones. This includes the network CIDR, gateway, reserved IP ranges, target logical switch, and vSphere cluster for the management and the Kubernetes availability zones. At least three Kubernetes availability zones are recommended for a highly available deployment of your chosen application runtime.

CANCEL **BEGIN**

- 4 Select the license agreements and click **Next**.

- 5 On the General Settings page, type a name for the Enterprise PKS solution.

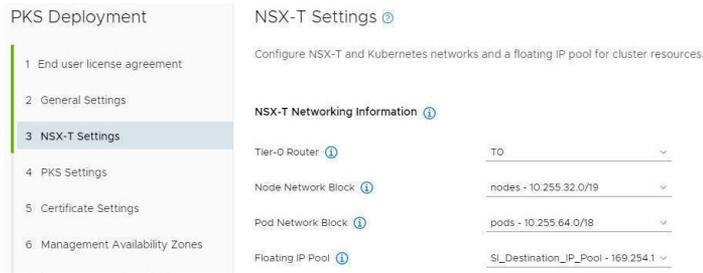
- 6 Select the NSX-T workload domain where you want to deploy Enterprise PKS.
- 7 Type a password for the Administrator user.
The password must contain nine or more characters and can be a mix of uppercase and lowercase characters, numbers, and special characters`
- 8 Re-type the Administrator password.
- 9 In the PKS Master section, type a user name for the Enterprise PKS API account.
The Enterprise PKS API account is used for performing cluster operations.
- 10 Type a password for the Enterprise PKS API account.
The password must contain nine or more characters and can be a mix of uppercase and lowercase characters, numbers, and special characters`
- 11 Re-type the password for the Enterprise PKS API account.
- 12 Enter a decryption password.
The decryption passphrase is used to encrypt the Operations Manager system. You must store this value securely as it needs to be entered after each Operations Manager reboot. The passphrase cannot be retrieved or reset.
The password must contain nine or more characters and can be a mix of uppercase and lowercase characters, numbers, and special characters.
- 13 Re-type the decryption password.
- 14 Click **Next**.

Specify NSX-T Settings

Provide NSX-T router information and floating IP pool for cluster resources.

Procedure

- 1 Select the tier-0 gateway (sfo01-w-tier-0-01 in our example).



- 2 Select the IP blocks and IP pool. Refer to the table below that includes values from our example.

Setting	Example Value
IP block for Nodes	sfo01-w-nodes-ip-block
IP block for Pods	sfo01-w-pods-ip-block
IP Pool	sfo01-w-floating-ip-pool

- 3 Click **Next**.

Specify PKS Settings

Provide the FQDN and network settings for Enterprise PKS API and Pivotal Operations Manager.

Procedure

- 1 Type the FQDN to access the Enterprise PKS API. For example, api.pks.example.com.
- 2 Specify whether you want to deploy Harbor Registry.
- 3 Select the datastore where the Enterprise PKS management components will be deployed.
- 4 Specify the following settings for PKS. For more information on the example values in the table, refer to [Prerequisites for Deploying VMware Enterprise PKS](#).

Setting	Example Value
FQDN	https://sfo01w02ops01.sfo01.rainpole.local
VM Name	sfo01w01ops01
IP Address	10.0.0.10
Gateway	10.0.0.1
Netmask	255.255.255.0
DNS	172.16.11.5
DNS Suffix	rainpole.local
NTP	172.16.11.251

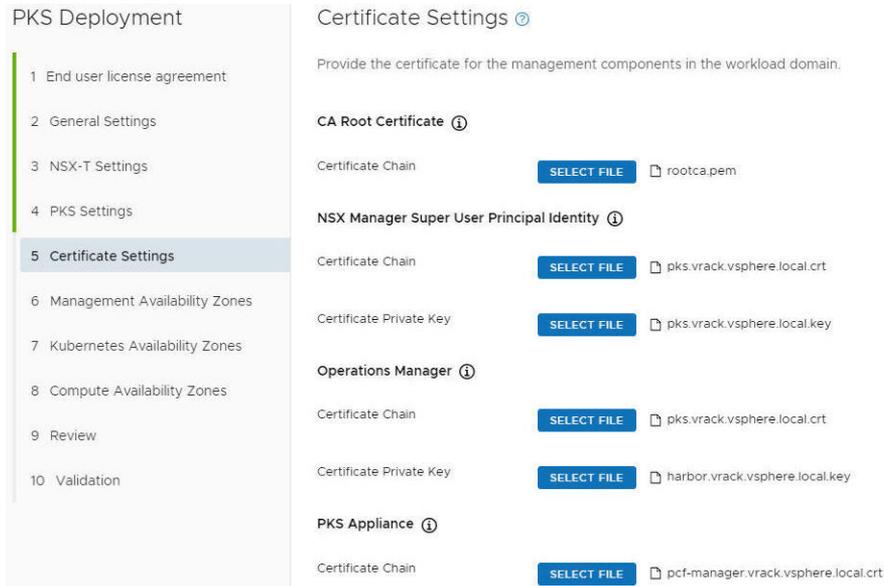
5 Click **Next**.

Provide Certificate Details

Provide certificate details for NSX Manager super user and Enterprise PKS management components. You had generated these certificates as part of the prerequisites for deploying Enterprise PKS.

Procedure

1 On the Certificates Settings page, select the certificate files for the listed components.



2 Click **Next**.

Specify Management Availability Zones

Specify information about Enterprise PKS management network and management availability zones.

Procedure

1 Specify the following settings for the Enterprise PKS management network . For more information on the example values in the table, refer to [Prerequisites for Deploying VMware Enterprise PKS](#).

Setting	Example Value
Management Network CIDR	10.0.0.0/24
Management Gateway	10.0.0.1
Management IP Reserved Range	10.0.0.1-10.0.0.10
Management DNS	172.16.11.5
Management Network	sfo01-w-pks-mgmt

PKS Deployment

- 1 End user license agreement
- 2 General Settings
- 3 NSX-T Settings
- 4 PKS Settings
- 5 Certificate Settings
- 6 Management Availability Zones**
- 7 Kubernetes Availability Zones
- 8 Compute Availability Zones
- 9 Review
- 10 Validation

Management Availability Zones ?

Provide the network settings and select an availability zone for the PKS management components.

Management Network ⓘ

Management portgroup ⓘ mgmt-ls

Management Network CIDR ⓘ 10.255.0.0/24

Management Gateway ⓘ 10.255.0.1

Management IP Reserved Range ⓘ 10.0.0.1 To 10.0.0.10

Management DNS ⓘ 172.16.11.5

Management Availability Zone ⓘ

Availability Zone Name ⓘ sfo01-w02-shared01

Cluster Name ⓘ clus3

Resource Pool ⓘ AZ1

- 2 Specify the following settings for the Enterprise PKS management availability zone . For more information on the example values in the table, refer to [Prerequisites for Deploying VMware Enterprise PKS](#).

Setting	Example Value
Management Availability Name	sfo01-w02-shared01
Cluster Name	cluster1
Resource pool in Cluster	RP-SahredAZ

- 3 Click Next.

Specify Kubernetes Availability Zones

Kubernetes availability zones are for legacy workloads. They are not used in the Enterprise PKS solution deployed by Cloud Foundation. However, you must complete the fields on this page to proceed.

Procedure

- 1 Specify the CIDR, gateway, reserved IP range, and DNS server from IP subnet that is routable to your corporate network and to the SDDC management network as the logical switch. The tier-1 router you created must have the gateway address set on the port connected to logical segment pks-service.

Setting	Example Value
Kubernetes Network CIDR	10.0.1.0/24
Kubernetes Gateway	10.0.1.1
Kubernetes IP Reserved Range	10.0.1.1 - 10.0.1.100

Setting	Example Value
Kubernetes DNS	172.16.11.5
Kubernetes Network	sfo01-w-pks-service

2 Click **Next**.

Configure Compute Availability Zones

Kubernetes cluster nodes are deployed in the compute availability zones.

Procedure

1 In the Add Availability Zone section on the Compute Availability Zones page, type a name for the availability zone.

Availability Zone Name	Cluster Name	Resource Pool
<input type="checkbox"/> sfo01-w02-comp01	clus3	AZ1
<input type="checkbox"/> sfo01-w02-comp02	clus3	AZ2
<input type="checkbox"/> sfo01-w02-comp03	clus3	AZ3

2 Select the cluster and resource pool for the availability zone and click **Add**. See the table below for values from our example.

3 Click **Add**.

The availability zone is listed in the Availability Zones Added section.

4 Repeat steps 1- 3 to add the remaining availability zones.

Refer to the table below for values from our example.

Name	Cluster	Resource Pool
sfo01-w02-comp01	cluster1	RP-Comp1AZ
sfo01-w02-comp02	cluster1	RP-Comp2AZ
sfo01-w02-comp03	cluster1	RP-Comp3AZ

5 Click **Next**.

Review Summary

Review the deployment summary.

Procedure

- 1 On the Review Summary page, review the deployment summary.

The screenshot shows the 'Validation' step in the PKs Deployment wizard. The left sidebar lists steps 1 through 10, with '10 Validation' selected. The main area displays a table of tasks being validated, with a 'Status' dropdown menu open. The table lists various tasks such as 'Semantic validation action', 'Validate Passwords', 'ProxyContractAction', etc., all with a status of 'SUCCESSFUL'. The 'Status' dropdown menu is open, showing options: ALL, NEW, INFO, IN PROGRESS, FAILED, ERROR, PENDING, and SUCCESSFUL. A 'DOWNLOAD' button is visible in the top right corner.

Tasks	Status
Semantic validation action	SUCCESSFUL
Validate Passwords	SUCCESSFUL
ProxyContractAction	SUCCESSFUL
Prepare PCF files validation action	SUCCESSFUL
PCF bundle validation action	SUCCESSFUL
Assemble validation output	SUCCESSFUL
Validate WLD is eligible for PKS installation action	SUCCESSFUL
Prepare PCF workload domain validation action	SUCCESSFUL
Validate resource pools are in cluster	SUCCESSFUL
Validate cluster storage	SUCCESSFUL
Validate cluster memory	SUCCESSFUL
Validate VM name uniqueness within the vCenter	SUCCESSFUL
Assemble validation output	SUCCESSFUL

To make an edit, click **Back**.

- 2 Click **Next**.

The input you provided in the wizard is validated and the list of tasks being validated is displayed on the Validation page. If a validation fails, click Retry. When the validation is successful, the Finish button is enabled.

- 3 Click **Finish**.

Results

As each task is validated, the status for that task changes to Successful. You can sort the task by status. You can also download the list of tasks being validated by clicking **Download**.

If all validation tasks are not successful, you can either ignore or retry them based on their severity.

After the validation is completed, the Enterprise PKS solution is added to the PKS table on the Workload Domains page. The Tasks table displays details of the tasks being performed and the status of each task.

View PKS Solution Details

The Enterprise PKS page displays high level information about the PKS solutions in the Cloud Foundation system. CPU, memory, and storage utilized by the workload domain is also displayed here.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domain**.
- 2 In the PKS section, click **View Details**.

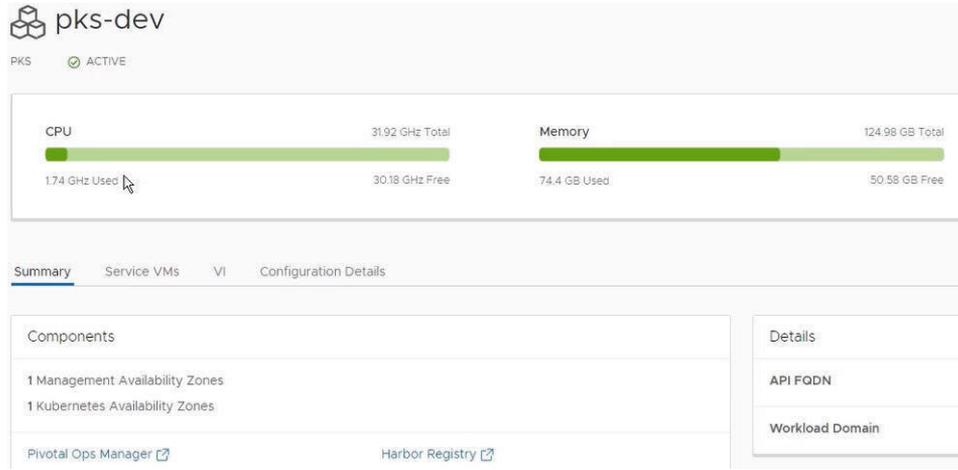
The PKS table displays a list of PKS solutions.

- 3 Click a solution name.

The solution details page displays CPU, memory, and storage allocated to the domain. The tabs on the page display additional information as described in the table below.

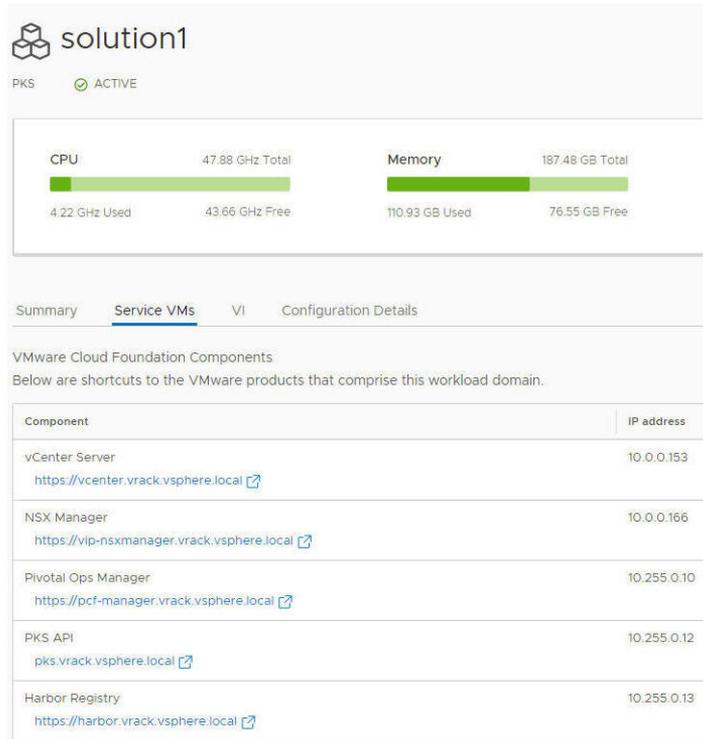
Tab Information Displayed

- Summary
- Solution components
 - Links to the Operations Manager and management portal for Harbor Registry.
 - PKS API FQDN for connecting to the PKS CLI tool
 - You can use the PKS CLI tool to create Kubernetes clusters.
 - NSX-T workload domain where this solution is deployed.



Service VMs

Direct links and IP addresses for the service VMs deployed for this Enterprise PKS solution.



Tab Information Displayed

VI Details about the NSX-T VI workload domain where this solution is deployed.

The screenshot shows the 'VI' tab for a 'solution1' workload domain. At the top, there is a 'DOWNLOAD CONFIG' button and a 'PKS ACTIVE' indicator. Below this, two progress bars show resource usage: CPU (47.88 GHz Total, 4.22 GHz Used, 43.66 GHz Free) and Memory (187.48 GB Total, 110.93 GB Used, 76.55 GB Free). A navigation bar includes 'Summary', 'Service VMs', 'VI' (selected), and 'Configuration Details'. A table below lists workload domains:

Domain	Type	CPU Usage	Memory Usage	vSAN Storage Usage	Status	Owner	Cluster	Hosts	Updates Available
nsxt-vi	VI	9%	59%	-	ACTIVE		2 clusters	6	Up-to-date

Configuration Details Solution , NSX-T, PKS, management, and Kubernetes settings provided when the solution was deployed.

The screenshot shows the 'Configuration Details' tab for a 'solution1' workload domain. It features the same resource usage bars as the previous screenshot. The navigation bar includes 'Summary', 'Service VMs', 'VI', and 'Configuration Details' (selected). The configuration settings are organized into sections:

- General Settings**
 - Workload Domain: nsxt-vi
- NSX-T Settings**
 - TO Router Id: e3c5227b-fdd0-482c-9523-2b5c9e565887
 - Node Network Block Id: 23e48d46-3557-48e2-904a-50a441ca8460
 - Floating IP Pool Id: d60764d6-9b13-4a80-98b1-e2287573ab2d
 - Pods IP Block: a8865c9c-095b-44f3-b8ee-fd0309ebacb9
- PKS Settings**
 - PKS API FQDN: pks.vrack.vsphere.local
 - Pivotal Ops Manager FQDN: pcf-manager.vrack.vsphere.local
 - Pivotal Ops Manager FQDN VM Name: pcf-manager

Delete a Enterprise PKS Solution

You can delete an Enterprise PKS solution from your Cloud Foundation deployment.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
- 2 Click **View Details** in the PKS section.

- 3 Click the dots next to the name of the solution you want to delete and choose **Delete Solution**.

A confirmation window appears.

- 4 Type the name of the solution you want to delete.
- 5 Click **Delete PKS Solution** to proceed.

The task panel displays the progress of the delete workflow. When the removal process is complete, the solution is removed from the PKS table.

Expand or Shrink an NSX-T Workload Domain associated with an PKS Solution

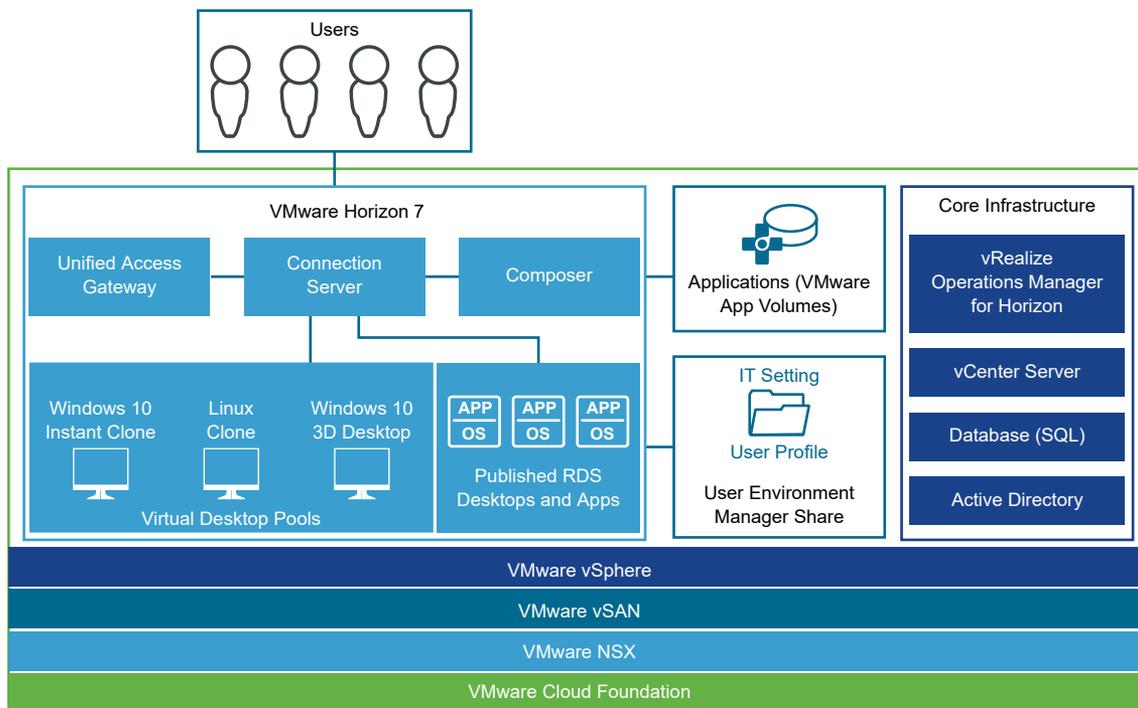
To expand an NSX-T workload domain associated with an Enterprise PKS solution, see [Expand a Workload Domain](#). You cannot shrink an NSX-T domain associated with an Enterprise PKS solution.

Working with Horizon Domains

10

A Horizon domain automates deployment of VMware Horizon components and supporting infrastructure to enable you to deliver Virtual Desktop Infrastructure (VDI) and Remote Desktop Session Host (RDSH) desktops and applications. These can be delivered as persistent, linked clone, or instant clone desktops. The Horizon domain can include VMware App Volumes for dynamic application mounting and User Environment Manager for a persistent end user experience.

Figure 10-1. Components of a Horizon Domain



The Horizon domain is decoupled from resource provisioning - one or more VI workload domains must be created before deploying a Horizon domain. During the domain deployment, one to three Connection Servers and a corresponding load balancer is deployed. In addition, you can choose the optional components that you want to deploy:

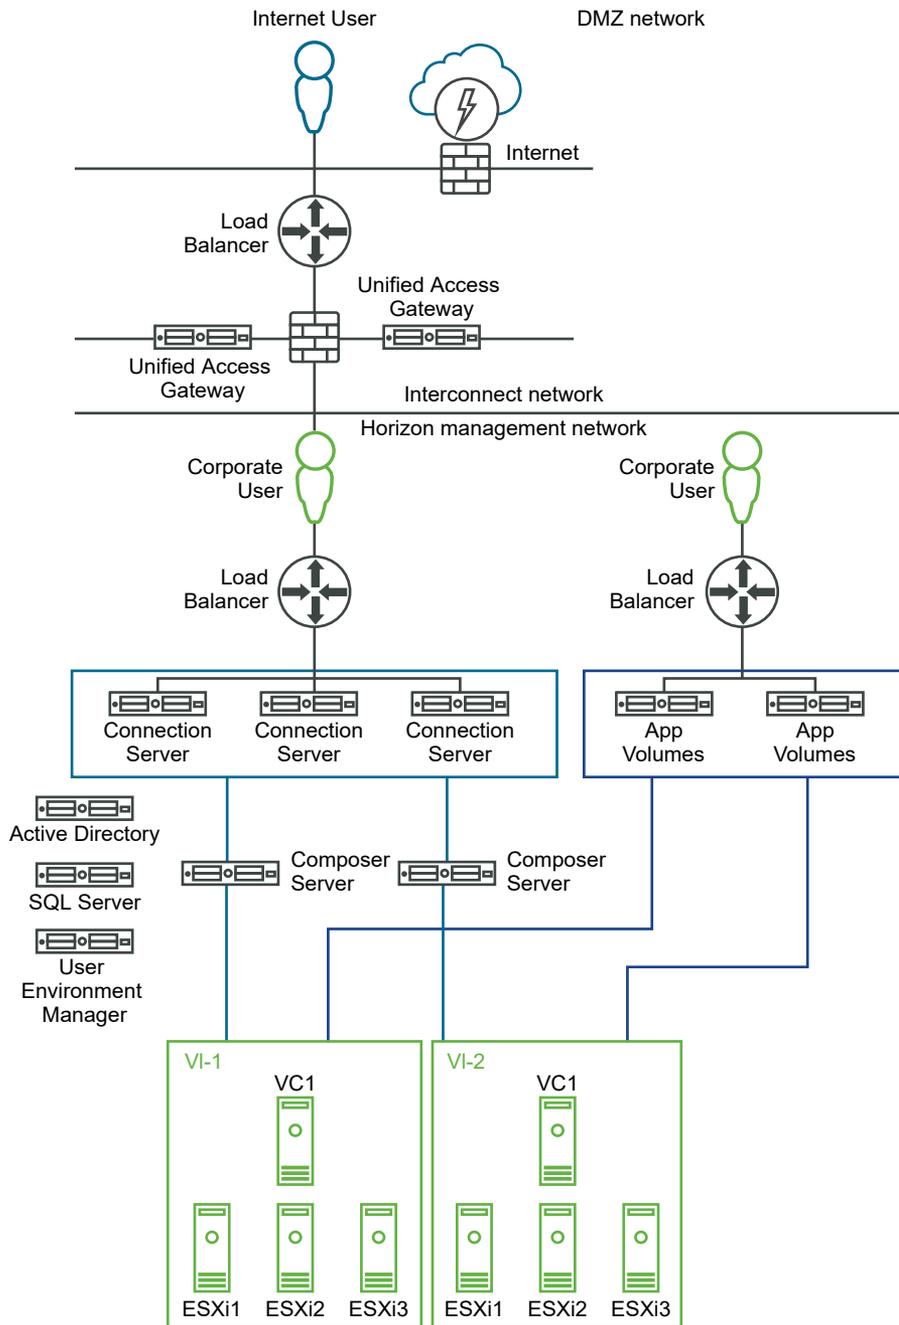
- Composer Server
- App Volumes

- User Environment Manager
- Unified Access Gateway

The architecture diagram below shows a Horizon domain deployed with the following components:

- three Connection Servers
- two VI workload domains
- two App Volumes Managers
- one User Environment Manager
- two Composer Servers
- two Unified Access Gateway appliances
- three Load Balancers (one for incoming WAN traffic across Unified Access Gateway, one across Connection Servers, and one across App Volumes Manager)

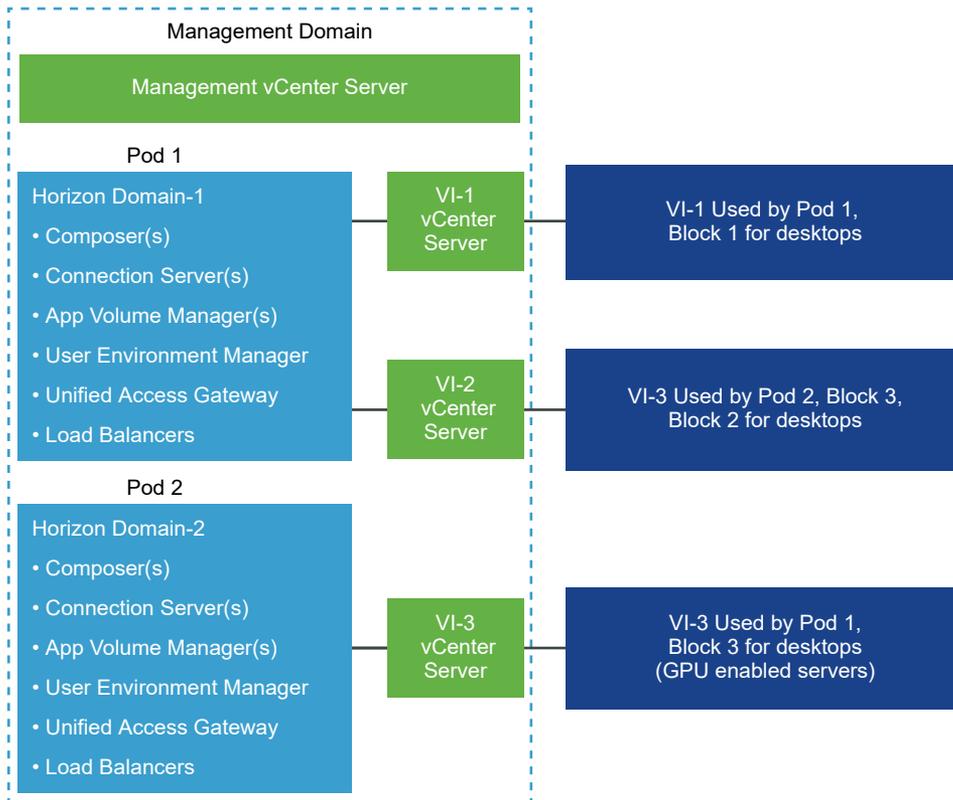
Figure 10-2. Deployed Horizon Domain



The Horizon domain is based on the Horizon Reference Architecture, which uses Pod Block architecture to enable you to scale as your use cases grow. For more information about the architecture and number of supported virtual machines, refer to the Horizon 7 Pod and Block section in the [VMware Workspace ONE and VMware Horizon 7 Enterprise Edition On-premises Reference Architecture](#) document. A Horizon domain maps to a single pod. To support multiple

Pods, you can deploy multiple Horizon domains (where each domain maps to a single pod). If Cloud Pod Architecture is required to connect the pods for advanced multi-site or scale out workflows, this can be done manually. See Cloud Pod Architecture Overview in the [Horizon 7 Architecture Planning document](#).

Figure 10-3. Horizon Domain Pod and Block



This chapter includes the following topics:

- [Sizing Guidelines](#)
- [Prerequisites for a Horizon Domain](#)
- [Create a Horizon Domain](#)
- [Resume Horizon Domain Creation](#)
- [Exporting and Importing a Horizon Domain Configuration](#)
- [View Horizon Domain Details](#)
- [Expand a Horizon Domain](#)
- [Delete Horizon Domain](#)

Sizing Guidelines

This section describes sizing considerations to be used when planning your Horizon domains.

Sizing VI Workload Domain Capacity for Use With Horizon Domains

Sizing of hosts for a Horizon domain is complicated. Typically, this involves a user study to ensure that the workloads being used by your enterprise are well understood. For example, heavy graphical use versus simple web based applications can make a big difference to requirements.

The Digital Workspace Designer can size and estimate the hardware required to run desktop or RDSH workloads and provide server components numbers and specifications. You can then work with VMware or a certified partner to validate the input and unlock the detailed results. Refer to the [Digital Workspace Tech Zone](#) to use the online sizing tool.

Sizing a Horizon Domain

The Horizon domain is based on the Horizon Reference Architecture, which uses the Pod and Block architecture to enable you to scale as your use cases grow. In Cloud Foundation, a Horizon domain is equivalent to a pod and a VI workload domain is equivalent to a block. Multiple VI workload domains can be associated with a single Horizon domain. This allows scale out to the recommended maximum of 10,000 desktops per Horizon domain.

A high level summary of scale considerations is as follows:

- A Horizon domain (pod) can deliver a recommended maximum of 10,000 desktops
- vCenter Server is the delimiter of a VI workload domain (block). The number of recommended VMs per vCenter Server (and therefore per block) depends on the types of VMs in use:
 - 5,000 instant clone VMs (without App Volumes)
 - 4,000 linked clone or full clone VMs (without App Volumes)
 - 2,000 VMs with App Volumes

Hence, for a Horizon domain to support the maximum recommended 10,000 desktops with App Volumes, you would need five VI workload domains.

Other sizing considerations are as follows:

- One Connection Server appliance per 2,000 Horizon server connections up to a maximum of seven servers
- One Unified Access Gateway appliance per 2,000 desktop connections
- One Composer Server per vCenter Server if traditional clones are used

It is recommended that you have two or more Connection Servers, Unified Access Gateway appliances, and App Volumes Manager instances to ensure high availability even if usage is less than the per-server maximum.

Cloud Foundation also supports multiple Horizon domains (multiple pods). However, Cloud Foundation does not provide automation for Cloud Pod Architecture to connect these pods. For information on manual steps, see Cloud Pod Architecture Overview in the [Horizon 7 Architecture Planning document](#). If Cloud Pod Architecture is required to connect the pods for advanced multi-site or scale out workflows, this can be done manually.

SQL Server Sizing

A Horizon domain uses the SQL Server for storing logs and relatively static data, so there are no heavy performance requirements. For enterprise deployments, the reference architecture recommends clustered databases to be used for redundancy, especially for App Volumes and Composer Servers databases.

You can use one or more SQL Server instances for a Horizon domain. For example, you can use multiple clustered instances to separate event data from runtime data and entitlement information.

Prerequisites for a Horizon Domain

Complete each prerequisite in this section before creating a Horizon domain.

Horizon 7 License

You must have a valid Horizon 7 license key, purchased separately from the Cloud Foundation license. You must add this license key to Cloud Foundation. See [Add License Keys for the Software in Your Cloud Foundation System](#).

Horizon 7 Install Bundle

Download the Horizon 7 install bundle. See [Chapter 13 Downloading an Install Bundle](#).

Networks

The following networks must be configured.

- DMZ network

The DMZ network is the intermediate network between the corporate network and the internet. The incoming interface of the Unified Access Gateway appliances and the DMZ load balancer are connected here.

- Interconnect network

This is an optional network for high security environments. The outgoing interface of the Unified Access Gateway appliances are connected with the management network here. This network must be routable to the Horizon management network.

Instead of having an interconnect network, you can also connect Unified Access Gateway appliances directly to the Horizon management network.

- Horizon management network

The Horizon management is the network dedicated to the Horizon components. All Horizon VMs (except Unified Access Gateway) must be on this network. All Connection Servers, Composer Servers, App Volumes, User Environment Manager and management interface of the Unified Access Gateway appliances must have IP addresses from this network. In addition, the load balancers deployed by Horizon domain in front of the Connection Servers and App Volumes must be in this network as well.

Unified Access Gateway has three interfaces - internal, external, and management:

- The internal interface can be either in the Horizon management or interconnect network. If it is on the interconnect network, it must be routable to the Horizon management network.
- External interface must be in the DMZ network.
- Management interface must be in the Horizon Management network.

Load Balancers and IP Addresses

External IP addresses must be available for all VMs and load balancers. The following components need load balancers:

- Connection Servers
- App Volumes (optional component)
- Unified Access Gateway appliances (optional component)

Load balancers must in the same network as the VMs they serve (Connection Servers and App Volumes load balancers in the Horizon management network and Unified Access Gateway load balancer in the DMZ network).

VXLAN Port Groups

VXLAN port groups must be created for the following:

- Horizon VMs in the Horizon management network
- Incoming interface (DMZ network)
- Outgoing interface (Interconnect network)

DNS Records

DNS records for load balancers must be pre-created such that the DNS names assigned to the entry points for load balancers are resolvable to the IP addresses being assigned to the load balancers. This is validated during the Horizon domain creation.

If Secure Dynamic Update is enabled within your Active Directory, a DNS record for each deployed Windows server is added automatically. If Secure Dynamic Update is disabled, you must create a DNS record for each Windows server you are planning to deploy. User Environment Manager, Connection Servers, App Volumes, and Composer Servers are Windows servers.

Custom Windows Image

You must provide a Windows Server image in OVA format for use with the Windows server components. This allows you to configure those server images according to your corporate guidelines. Cloud Foundation supports Windows 2016 and Windows 2012r2 images with the latest VMware tools installed and Windows Remote Management (WinRM) enabled. The template must have an administrator user account enabled.

Active Directory

You need two groups in your Active Directory for a Horizon domain. During the Horizon domain creation, one group is assigned administrative privileges for the Connection Servers and the other group is assigned administrative privileges for App Volumes. You can use a single group with privileges for both. Note that you cannot use groups of Builtin Local type.

You also need two service accounts in your Active Directory. The first account is required for Composer Servers. You can either have a dedicated account for each Composer Server, or one account for all Composer Servers. The second service account must have read-write permissions for the Organizational Unit. This account is used to join the servers that are deployed by Horizon.

A Horizon administrator account is also required for logging in to Horizon and App Volumes. This user must be a member of both the Horizon and App Volumes groups.

All users must be added with the following syntax:

domainName\username

where the domain name is the FQDN of the domain and user name matches the user logon name in AD Users and Computers console (pre-Windows 2000). For example `horizon-1.local\vdiaadmin`.

SQL Servers

You may either use one SQL Server for your entire environment, or use one SQL Server per deployed component. A user account with permissions to create databases is required for each SQL Server to be used. One account can be used for all SQL servers, or you can have a dedicated account per server. Each user account must be an SQL user.

Connection Servers, Composer Servers, and App Volumes require SQL databases. A dedicated SQL database is required for each Composer Server. All Connection Servers share one database, and all App Volumes can share one database. As an example, if you have five Composer Servers and an App Volumes in your environment, you will need seven SQL databases - five for the Composer Servers and one each for App Volumes and Connection Servers. All seven SQL databases can be inside a single SQL Server instance.

VI Workload Domains

You must pre-create the required VI workload domains, which are then associated with the Horizon domain. The end user desktops are placed on the VI workload domains.

Create a Horizon Domain

You configure a Horizon domain through the Horizon domain wizard. The databases are created during the domain creation process. If a database with the specified name exists, it is overwritten. The information you provide in the wizard is automatically saved when you proceed to a new page in the wizard so that you can resume configuration if you leave the wizard before completing the configuration.

Prerequisites

Ensure that you have completed all prerequisites described in [Prerequisites for a Horizon Domain](#).

Procedure

- 1 On the SDDC Manager Dashboard, click **+ Workload Domain** and then click **Horizon**.
- 2 On the Horizon Domain Creation page, select **Create a new Horizon Domain**.

Horizon Workload Domain Creation ⓘ ×

Create a new Horizon Domain or resume previous session.

Create a new Horizon Domain

Resume Configuration

Last edited at 05:45:17 03/06/2019

Upload Configuration

CANCEL NEXT

- 3 Click **Next**.
- 4 On the Horizon Domain Configuration Checklist page, confirm that you have met all listed prerequisites by selecting the check boxes.
- 5 Click **Start Configuration Wizard**.
- 6 On the General page, enter the general parameters for the Horizon domain.

Field Name	Information to be Entered
Domain Name	Enter a name for the Horizon domain.
Domain VM Name Prefix	Enter a prefix for the management virtual machines created for this domain. This can be overridden later.

Field Name	Information to be Entered
Windows File	<p>Select whether you want to upload a Windows image for the Windows based server components, or use an existing OVA template.</p> <p>To use an existing OVA template, the Windows image must have been uploaded to the management vCenter Server. In addition:</p> <ul style="list-style-type: none"> ■ VM name on vCenter Server where the OVA template was imported must start with vdi-vm-. ■ VM must be powered off.
Windows OVA Template	<p>This field appears if you selected Upload OVA Template in the Windows File field. Click Upload and browse to select the file. This image is used as the operating system for the Horizon management components.</p> <p>The upload progress is displayed. Since the image is a large file, the upload may take several minutes. You cannot proceed till the image is uploaded.</p>
Windows OVA Template Path	<p>This field appears if you selected Use Existing OVA Template in the Windows File field. Select the template to be used.</p>
Admin Username	<p>Enter the administrator user name for the Windows image you uploaded. This user name must exist in the image you uploaded and be enabled. This user must always be the administrator.</p>
OVA Windows Administrator Password	<p>Enter a password for the administrator user on the VMs to be deployed using this template.</p>
Confirm Password	<p>Re-enter the administrator password.</p>
Management Port Group	<p>Select the management VXLAN port group on which the Horizon management components will be connected. You must have pre-created this port group for it to appear in this drop-down.</p>

7 Click **Next**.

Select VI Workload Domains for the Horizon Domain

A Horizon domain can deliver a recommended maximum of 10,000 desktops. You must select one or more pre-created VI workload domains where the desktops are to be placed. A VI workload domain can accommodate 5,000 desktops with traditional clones, 4,000 desktops with traditional clones, and 2,000 desktops with App Volumes.

Procedure

- 1 On the Select Existing VI Domains page, select the VI workload domains for the Horizon domain.
- 2 Click **Next**.

Provide Active Directory Details

Provide details for the Active Directory to which you want to connect the Horizon domain.

Procedure

- 1 Provide the following information for the Active Directory.

Field Name	Information to be Entered
FQDN	Enter the Fully Qualified Domain Name (FQDN) for the Active Directory.
Organizational Unit for Horizon VMs	Enter the Organizational Unit where Horizon management VM are to be placed. The OU format must be as follows: OU=aaa, DC=bbb, DC=ccc For example, ou=vdi, dc=horizon-1, dc=local
Administrator Username	Enter a user name from the Active Directory with administrative privileges granted for the Horizon domain. The user must be a member to the Active Directory groups managing Horizon and App Volumes. This user is only used during the creation process - it can later be deleted or the password can be changed. Format must be as follows: <i>myDomain.local\myUserName</i> For example, horizon-1.local\vdiadmin
Administrator Password	Enter the password for the specified Administrator user.
Read Write Account	Enter an Active Directory account name with read and write permissions on the Active Directory. The user must have read and write permissions for the specified Organization Unit, and will join the servers under its context. Format must be as follows: <i>myDomain.local\myUserName</i> For example, horizon-1.local\rwadmin
Read Write Account Password	Enter the password for the specified read and write account.
LDAPS	Select this to use Secure LDAP (LDAPS) to connect to the Active Directory.
DC 1 IP Address	Enter the IP address for Domain Controller 1.
Thumbprint	If you are using LDAPS and are configuring Domain Controller 1, validate the thumbprint of Domain Controller 1.
DC 2 IP Address	Enter the IP address for Domain Controller 2. This is optional, but recommended.
Thumbprint	If you are using LDAPS and are configuring Domain Controller 2, validate the thumbprint of Domain Controller 2.

- 2 Click **Next**.

Provide SQL Server Details

You can one or more SQL Server instance for the Horizon domain. For example, you can use a different SQL Server for events and runtime data.

Procedure

1 In the Add SQL Servers section, select **Add manually** to add SQL Servers information manually or **Import from JSON template** to import SQL Servers from a JSON file.

- To add SQL Servers manually, follow the steps below.
 - a Provide the following information.

Field Name	Information to be Entered
Alias	Enter an alias for the SQL Server.
FQDN	Enter the FQDN for the SQL Server.
SQL Instance Name	Enter name of the the SQL Server instance to be used for the Horizon domain.
SQL Port	Enter the port number to connect to the SQL Server. The default port is 1433.
Database Username	Enter the user name to connect to the SQL Server.
Database Password	Enter the password for the database user name.
Confirm Database Password	Re-enter the password for the database user name.

b Click **Add**.

The SQL Servers is added to the SQL Servers table.

c Repeat steps a and b for additional SQL Servers as required.

- To import SQL Servers from a JSON file, click **Browse**, select the file, and click **Upload**.
SQL Servers from the JSON file are added to the SQL Servers table.

2 Click **Next**.

Add Load Balancers

Depending on the Horizon domain components being deployed, you must configure one to three load balancers.

Load balancers are required for the following.

- Load balance incoming internal requests (and Unified Access Gateway appliance south bound traffic) across the Connection Servers. This load balancer is mandatory.
- Load balance incoming WAN based requests across the Unified Access Gateway appliances. This load balancer is required only if you are deploying an Unified Access Gateway appliance.
- Load balance desktop connect requests across App Volumes Managers. This load balancer is required only if you are deploying App Volumes.

See Fig 2 in [Chapter 10 Working with Horizon Domains](#) for an example.

Procedure

- 1 On the Load Balancers page, provide a prefix for the load balancer VM names.
- 2 In the Load Balancers section, select **Add manually** to add load balancers manually or **Import from JSON template** to import load balancers from a JSON file.
 - To add load balancer manually, follow the steps below

- a Provide the following information.

Field Name	Information to be Entered
Alias	Enter an alias for the load balancer.
FQDN	Enter the FQDN for the load balancer.
VM Name	Enter a name for the load balancer VM.
IP	Enter the IP address for the load balancer.
Subnet Mask	Enter the subnet mask for the load balancer.
Gateway	Enter the gateway for the load balancer.
CLI Password	Enter the CLI password for the load balancer. Specified password must meet the following guidelines: <ul style="list-style-type: none"> ■ subsequent identical characters. ■ Contain at least 12 characters and no more than 255 characters. ■ Start with an alphabetical character. ■ Contain at least one lowercase character. ■ Contain at least one uppercase character. ■ Contain at least one digit. ■ Contain at least one special character. ■ Should not contain any whitespace.
Confirm CLI Password	Re-enter the CLI password for the load balancer.

- b Click **Add**.

The load balancer is added to the Load Balancer table.

- c Repeat steps a and b for additional load balancers as required.

- To import load balancers from a JSON file, click **Browse**, select the file, and click **Upload**.
Load balancers from the JSON file are added to the Load Balancer table.

- 3 Click **Next**.

Add Connection Servers

Connection Servers are mandatory for a Horizon domain. It is recommended that you have a minimum of two Connection Servers for high availability; the maximum number supported by a Horizon domain is seven. A Connection Server can scale to a maximum of 2000 active sessions. You can reach the maximum recommended 10,000 sessions with five Connection Servers but

two servers are required for high availability. Hence, seven Connection Servers are recommended for the 10,000 sessions.

Procedure

- 1 On the Horizon Connection Servers page, provide general information about the Connection Server.

Table 10-1. Horizon Connection Servers General

Field Name	Information to be Entered
Horizon License	Select the Horizon license key to use for this Horizon domain.
VM Name Prefix	Enter a prefix for the Connection Servers VM names.
Admin Group Name	Enter the administrator group name that has administrator access to the Horizon environment. This group must exist in the Active Directory. For example, horizonAdmin.
Load Balancer Alias	Select the load balancer to for the Horizon Connection Servers. The selected load balancer must be in the same network as the Connection Servers.
SQL Server Alias	Select the SQL Server to use with the Horizon Connection Servers.
Database Name	Enter the database name to use with Horizon Connection Servers. If a database with this name exists, it will be overwritten.

- 2 In the Add Connection Servers section, select **Add manually** to add Connection Servers manually or **Import from JSON template** to import Connection Servers from a JSON file.
 - To add the Connection Servers manually, follow the steps below

- a Provide the following information.

Field Name	Information to be Entered
Computer Name	Enter a computer name for the Connection Server.
FQDN	The FQDN for the Connection Server is auto-generated using the computer name and the Active Directory FQDN.
VM Name	Enter a Virtual Machine name for the Connection Server.
IP Address	Enter an IP address for the Connection Server.
Subnet Mask	Enter the subnet mask for the Connection Server.
Gateway	Enter a gateway for the Connection Server.

- b Click **Add**.

The Connection Server is added to the Connection Servers table.

- c Repeat steps a and b for additional Connection Servers as required.

- To import Connection Servers from a JSON file, click **Browse**, select the file, and click **Upload**.

Connection Servers from the JSON file are added to the Connection Servers table.

- 3 Click **Next**.

Add Composer Servers

A Composer Server is an optional component.

A Composer Server manages traditional clones using linked-clone technology. Each Composer server is paired with a vCenter Server. A block architecture with one vCenter Server per 4,000 linked clone VMs would require one Composer server. High availability is provided by VMware vSphere High Availability (HA), which restarts the Composer VM in the case of a vSphere host outage.

Composer Servers are not required if you plan to use only instant clones.

Procedure

- 1 Slide the Deploy Composer Servers toggle option to green.
- 2 In the Add Composer Servers section, select **Add manually** to add Composer Servers manually or **Import from JSON template** to import Composer Servers from a JSON file.
 - To add Composer Servers manually, follow the steps below.

- a Provide the following information.

Field Name	Information to be Entered
VM Name Prefix	Enter a prefix for the Composer Server VMs.
IP Address	Enter an IP address for the Composer Server.
Subnet Mask	Enter the subnet mask for the Composer Server.
Gateway	Enter a gateway for the Composer Server.
Computer Name	Enter a computer name for the Composer Server.
FQDN	The FQDN for the Composer Server is auto-generated using the computer name and the Active Directory FQDN.
VM Name	Enter a name the Composer Server VM.
Composer Service Account	Enter the user name for the Composer Service Account. For example, ComposerAdmin.
Composer Service Password	Enter the password for the Composer Service Account.
SQL Server Alias	Select the SQL Server to use with the Composer Server.
Database Name	Enter a name for the Composer Server database. If a database with the specified name exists, it will be overwritten.
Manage vCenter	Select the vCenter Server to pair with the Composer Server.

- b Click **Add**.

The Composer Server is added to the Composer Servers table.

- c Repeat steps a and b for additional Composer Servers as required.

- To import Composer Servers from a JSON file, click **Browse**, select the file, and click **Upload**.

Composer Servers from the JSON file are added to the Composer Servers table.

- 3 Click **Next**.

Add Unified Access Gateway Appliances

The Unified Access Gateway appliance is an optional component. If you do not deploy this while creating a domain, you can add it later by expanding the domain.

Unified Access Gateway appliances provide a secure means to allow WAN based user traffic to connect to Horizon desktops running in a Cloud Foundation datacenter. Unified Access Gateway appliances are also used for HTML access on Horizon desktops (browser based connectivity to desktops and applications).

Maximum active connections recommended per Unified Access Gateway appliance is 2000.

Procedure

- 1 Slide the Deploy Unified Access Gateway toggle option to green.
- 2 On the Horizon Unified Access Gateway page, provide general information about the appliance.

Table 10-2. Unified Access Gateway General

Field Name	Information to be Entered
VM Name Prefix	Enter a prefix for the Unified Access Gateway appliance VM names.
Load Balancer Alias	Select the load balancer to use with the Unified Access Gateway appliance. This must be a different load balancer than the one used for Connection Servers or App Volumes.
External (DMZ) Port Group	Select the DMZ port group.
Internal Port Group	Select the internal port group.
Management Port Group	The management port group you selected on the General tab of the wizard is displayed here. You cannot change this here.

- 3 In the Add Unified Access Gateway Appliances section, select **Add manually** to add Unified Access Gateway appliances manually or **Import from JSON template** to import Unified Access Gateway appliances from a JSON file.
 - To add Unified Access Gateway appliances manually, follow the steps below
 - a Provide the following information.

Field Name	Information to be Entered
VM Name	Enter a name for the Unified Access Gateway appliance VM.
Default Gateway	Enter the default gateway for the Unified Access Gateway appliance.
Admin Password	Enter a password for the Unified Access Gateway appliance.
External IP Address	Enter the external facing IP address for the Unified Access Gateway appliance.
External Subnet Mask	Enter the external subnet mask for the Unified Access Gateway appliance.

Field Name	Information to be Entered
Internal IP Address	Enter the internal IP address for the Unified Access Gateway appliance.
Internal Subnet Mask	Enter the internal subnet mask for the Unified Access Gateway appliance.
Management IP Address	Enter the management IP address for the Unified Access Gateway appliance.
External Subnet Mask	Enter the management subnet mask for the Unified Access Gateway appliance.

- b Click **Add**.

The Unified Access Gateway appliance is added to the Unified Access Gateway appliance table.

- c Repeat steps a and b for additional Unified Access Gateway appliances as required.

- To import Unified Access Gateway appliances from a JSON file, click **Browse**, select the file, and click **Upload**.

Unified Access Gateway appliances from the JSON file are added to the Connection Servers table.

- 4 Click **Next**.

Add App Volumes

App Volumes is an optional component. If you do not deploy this while creating a domain, you can add it later by expanding the domain.

App Volumes supports dynamic attachment of applications (AppStacks) to Horizon desktops based on user entitlement. It is recommended that you add two App Volumes for each VI workload domain associated with the Horizon domain.

Procedure

- 1 Slide the Deploy App Volumes toggle option to green.
- 2 In the App Volumes section, provide general information about the appliance.

Field Name	Information to be Entered
VM Name Prefix	Enter a prefix for the App Volumes appliance VM names.
Load Balancer Alias	Select the load balancer to use with the App Volumes. The load balancer must be in the same network as the App Volumes servers.
Admin Group	Enter the Active App Volumes Directory group that will be used to allow administrative access to the App Volumes management console. For example, AppVolAdmins. This does not have to be unique - it can be the same group as the one that administers the Connection Servers.
SQL Connection Alias	Select the SQL Server to use with App Volumes.
Database Name	Enter a database name for App Volumes. If a database with this name exists, it will be overwritten

- 3 In the Add App Volumes section, select **Add manually** to add App Volumes manually or **Import from JSON template** to import App Volumes from a JSON file.

- To add App Volumes Unified Access Gateway manually, follow the steps below

- a Provide the following information.

Table 10-3.

Field Name	Information to be Entered
IP Address	Enter the App Volumes IP address.
Subnet Mask	Enter the subnet mask for App Volumes.
Default Gateway	Enter the default gateway for App Volumes.
Computer Name	Enter a computer name for the App Volumes.
FQDN	The FQDN for App Volumes is auto-generated using the computer name and the Active Directory FQDN.
VM Name	Enter a name for the App Volumes VMs.

- b Click **Add**.

App Volumes is added to the App Volumes table.

- c Repeat steps a and b for additional App Volumes as required.

- To import App Volumes from a JSON file, click **Browse**, select the file, and click **Upload**.

App Volumes from the JSON file are added to the App Volumes table.

- 4 Click **Next**.

Add User Environment Manager

User Environment Manager is an optional component. If you do not deploy this while creating a domain, you can add it later by expanding the domain.

User Environment Manager enables per-user customization for desktops.

Procedure

- 1 Slide the Deploy User Environment Manager toggle option to green.

- 2 Provide the following information.

Field Name	Information to be Entered
VM Name Prefix	Enter a prefix for the User Environment Manager VM names.
IP Address	Enter an IP address for the User Environment Manager.
Subnet Mask	Enter the subnet mask for the User Environment Manager.
Gateway	Enter a gateway for the User Environment Manager.
FQDN	Enter the FQDN for User Environment Manager.
Computer Name	Enter a computer name for the User Environment Manager.

Field Name	Information to be Entered
VM Name	Enter a name the User Environment Manager VM.
Profile Folder	Enter the folder path for the home share of the profile.
Profile Sharename	Enter the profile sharename to use with User Environment Manager.
Configuration Folder	Enter the full path of the configuration folder to be used with User Environment Manager.
Configuration Sharename	Enter the configuration fileshare name.
Data Drive Size	Enter the required data drive disk size.

3 Click **Next**.

Review Horizon Domain Configuration

The Review page displays the Horizon domain configuration details.

Procedure

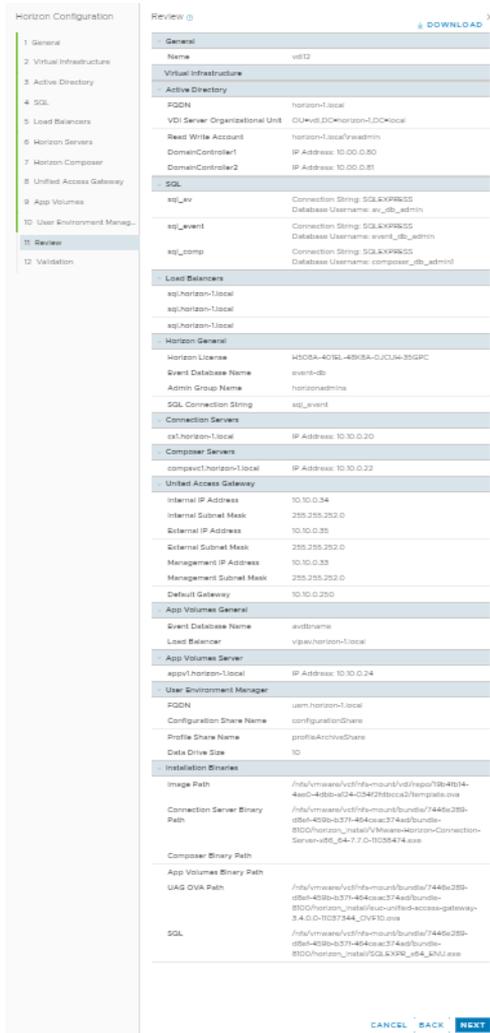
1 Review the domain configuration. Click **Back** to make edits to any section.

2 Click **Download** to download the domain configuration as a JSON file.

You can upload this JSON file to create a new Horizon domain.

3 Click **Next**.

The data you provided (credentials, FQDNs, portgroups, permissions, networking, etc.) is verified.



If there are any errors, you can resolve them and click **Retry** to run the validation again. For additional troubleshooting, review the log file on the SDDC Manager VM at `/var/log/vmware/vcf/solutionmanager/solutionmanager.log`.

Results

After the validation is complete, the Horizon domain is added to the Horizon table on the Workload Domains page. The Tasks table displays details of the tasks being performed and the status of each task. The domain creation takes some time to complete depending on the number of components being deployed.

Resume Horizon Domain Creation

The Horizon Domain wizard saves the information you entered. If you exited the wizard before completing the domain configuration, you can resume configuration from the time you left the wizard.

Procedure

- 1 On the SDDC Manager Dashboard, click **+ Workload Domain** and then click **Horizon**.
- 2 On the Horizon Domain Creation page, click **Resume Configuration**.

The date and time that you last edited the configuration is displayed.

Horizon Workload Domain Creation ? ×

Create a new Horizon Domain or resume previous session.

Create a new Horizon Domain

Resume Configuration

Last edited at 05:45:17 03/06/2019

Upload Configuration

CANCEL NEXT

- 3 Click **Next**.
- 4 On the Horizon Domain Configuration Checklist page, confirm that you have met all listed prerequisites by selecting the check boxes.
- 5 Click **Next**.
- 6 The General page displays the information you had entered earlier, but you must upload the Windows image again.
- 7 Click **Next** till you reach the page where you had left the wizard and then complete entering the remaining information. For more information on the required information, see [Create a Horizon Domain](#).

Exporting and Importing a Horizon Domain Configuration

You can export the configuration of a Horizon domain as a JSON file. You can use this file to create a similar Horizon domain configuration. The modified configuration file can be imported in Cloud Foundation to create a new Horizon domain.

Export a Horizon Domain Configuration

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domain**.
- 2 In the Horizon area, click **View Details**.

- 3 Click  next to the Horizon domain whose configuration you want to export and click **Download Config**.

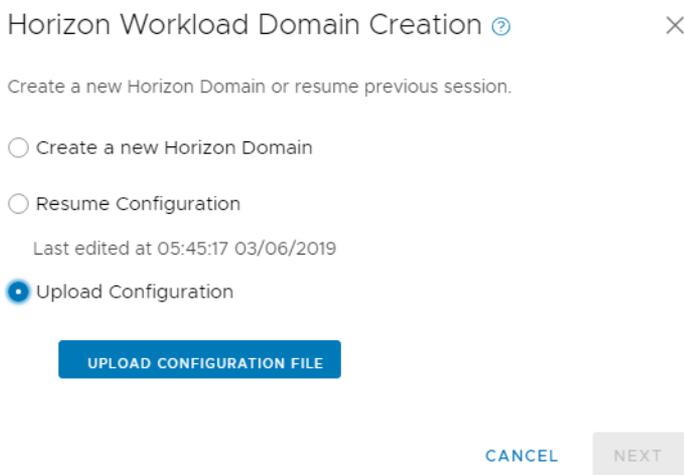
The configuration is downloaded as a JSON file. You can modify the JSON file to provide unique values for fields such as FQDNs, names, IP addresses, etc.

Import a Horizon Domain Configuration

You can import a configuration JSON file to create a new Horizon domain.

Procedure

- 1 On the SDDC Manager Dashboard, click **+ Workload Domain** and then click **Horizon**.
- 2 On the Horizon Domain Creation page, select **Upload Configuration**.



Horizon Workload Domain Creation  

Create a new Horizon Domain or resume previous session.

Create a new Horizon Domain

Resume Configuration

Last edited at 05:45:17 03/06/2019

Upload Configuration

UPLOAD CONFIGURATION FILE

CANCEL **NEXT**

- 3 Click **Upload Configuration File**.
- 4 Select the JSON file to be uploaded and click **Open**.
- 5 Click **Next**.
- 6 Progress through the domain creation wizard, making edits as required. For details on required information, see [Create a Horizon Domain](#).

Example: Sample Configuration JSON File

```
{
  "id": "1",
  "name": "vdi",
  "managementVcenters": [
    {
      "managementPortgroup": "VDI-DPortGroup-Mgmt",
      "dmzPortgroup": "dmz",
      "interconnectPortgroup": "VDI-DPortGroup-interconnect",
      "uagManagementPortgroup": "VDI-DPortGroup-Mgmt",
      "clusterName": "SDDC-Cluster1",
      "datastoreName": "sfo01-m01-vsan",
    }
  ]
}
```

```

    "datacenterName": "SDDC-Datacenter",
    "nsx": {
      "host": "nsxManager.vrack.vsphere.local",
      "password": "VMware1!"
    },
    "username": "administrator@vsphere.local",
    "password": "VMware123!",
    "host": "vcenter-1.vrack.vsphere.local",
    "psc": {
      "host": "psc-1.vrack.vsphere.local"
    }
  }
],
"resourceVcenters": [
  {
    "state": "DEPLOYED",
    "datacenters": [
      {
        "name": "new-vi-DC",
        "datastores": [
          {
            "name": "new-vi-vcenter-2-via-cluster1-vsan-01",
            "id": "sddc-ds-1",
            "hostIds": [
              "c7945622-d8c1-494e-aa76-180a720c1606",
              "801d55ec-156a-4b54-94fa-5019cd5a7d83",
              "203866dc-4ef3-42a8-880b-9d20656aa7de"
            ]
          }
        ]
      }
    ],
    "clusters": [
      {
        "name": "via-cluster1",
        "hosts": [
          {
            "name": "esxi-7.vrack.vsphere.local",
            "id": "c7945622-d8c1-494e-aa76-180a720c1606",
            "username": "root",
            "password": "EvoSddc!2016"
          },
          {
            "name": "esxi-6.vrack.vsphere.local",
            "id": "801d55ec-156a-4b54-94fa-5019cd5a7d83",
            "username": "root",
            "password": "EvoSddc!2016"
          },
          {
            "name": "esxi-5.vrack.vsphere.local",
            "id": "203866dc-4ef3-42a8-880b-9d20656aa7de",
            "username": "root",
            "password": "EvoSddc!2016"
          }
        ]
      }
    ]
  }
]

```

```

    }
  ],
  "username": "administrator@vsphere.local",
  "password": "VMware123!",
  "host": "vcenter-2.vrack.vsphere.local",
  "psc": {
    "host": "psc-1.vrack.vsphere.local"
  }
}
],
"peripheralServices": {
  "loadBalancers": [
    {
      "ipAddress": "10.10.0.25",
      "certificatePath": "/tmp/AppVCertificate7006187265877059244.p12",
      "certificatePassword": "LoEpxXC4",
      "fqdn": "lb1.horizon-1.local",
      "state": "DEPLOYED",
      "deployDetails": {
        "vmName": "lb1",
        "subnetMask": "255.255.255.0",
        "gateway": "10.10.0.250",
        "portgroup": "VDI-DPortGroup-Mgmt"
      }
    },
    {
      "ipAddress": "10.10.0.26",
      "certificatePath": "/tmp/ConServCertificate8496572034532614986.p12",
      "certificatePassword": "leHOTEZ9",
      "fqdn": "lb2.horizon-1.local",
      "state": "DEPLOYED",
      "deployDetails": {
        "vmName": "lb2",
        "subnetMask": "255.255.255.0",
        "gateway": "10.10.0.250",
        "portgroup": "VDI-DPortGroup-Mgmt"
      }
    },
    {
      "ipAddress": "10.20.0.27",
      "certificatePath": "/tmp/UagCertificate1734475760559824715.p12",
      "certificatePassword": "9EZhzEm0",
      "fqdn": "lb3.horizon-1.local",
      "state": "DEPLOYED",
      "deployDetails": {
        "vmName": "lb3",
        "subnetMask": "255.255.255.0",
        "gateway": "10.20.0.250",
        "portgroup": "dmz"
      }
    }
  ]
},
"activeDirectory": {
  "fqdn": "horizon-1.local",
  "netBiosName": "horizon-1",

```

```

    "vdiAdminUsername":"horizon-1.local\vdiadmin",
    "vdiAdminPassword":"VMware123!",
    "vdiServerOu":"OU\u003dvdi,DC\u003dhorizon-1,DC\u003dlocal",
    "ouRwUsername":"horizon-1.local\rwadmin",
    "ouRwPassword":"VMware123!",
    "securedAD":false,
    "domainControllers":[
      {
        "fqdn":"dc1.horizon-1.local",
        "ipAddress":"10.10.0.80"
      },
      {
        "fqdn":"dc2.horizon-1.local",
        "ipAddress":"10.10.0.81"
      }
    ]
  },
  "installDetails":{
    "imagePath":"/home/vcf/template.ova",
    "connectionServerBinaryPath":"/home/vcf/VMware-Horizon-Connection-Server-
x86_64-7.7.0-11038474.exe",
    "composerServerBinaryPath":"/home/vcf/VMware-viewcomposer-7.7.0-11038293.exe",
    "appVolumesServerBinaryPath":"/home/vcf/App Volumes Manager.msi",
    "uagOvaPath":"/home/vcf/euc-unified-access-gateway-3.4.0-11037344_OVF10.ova",
    "sqlExpressPath":"/home/vcf/SQLEXPRESS_x64_ENU.exe",
    "uemBinaryPath":"/home/vcf/VMware User Environment Manager 9.6.0.855 x64.msi",
    "connectionServerVersion":"7.5.5",
    "composerServerVersion":"4.5",
    "appVolumesServerVersion":"2.14",
    "uagOvaVersion":"buf",
    "sqlExpressVersion":"2017",
    "uemVersion":"uem",
    "ova":{
      "administratorUsername":"administrator",
      "administratorPassword":"VMware123!"
    },
    "updateKb2919442BinaryPath":"/home/vcf/download/Windows8.1-KB2919442-x64.msu",
    "updateKb2919355BinaryPath":"/home/vcf/download/Windows8.1-KB2919355-x64.msu",
    "dotNet462BinaryPath":"/home/vcf/download/NDP462-KB3151800-x86-x64-ALL0S-ENU.exe"
  },
  "sqlConnections":[
    {
      "fqdn":"sql.horizon-1.local",
      "id":"sqlext",
      "sqlInstanceName":"SQLEXPRESS",
      "sqlPort":1433,
      "state":"DEPLOYED",
      "dbUsername":"av_db_admin",
      "dbPassword":"VMware123!"
    }
  ]
},
"horizon":{
  "pods":[
    {

```

```

"eventDbName":"event-db-ext",
"sqlId":"sqlxt",
"vcFqdns":[
  "vcenter-2.vrack.vsphere.local"
],
"uagAppliances":[
  {
    "vmName":"uag1",
    "mgmtSubnetMask":"255.255.255.0",
    "mgmtIpAddress":"10.10.0.33",
    "externalSubnetMask":"255.255.255.0",
    "internalSubnetMask":"255.255.255.0",
    "defaultGateway":"10.0.0.253",
    "internalIpAddress":"10.10.0.34",
    "externalIpAddress":"10.20.0.35",
    "administratorPassword":"VMware123!",
    "state":"DEPLOYED"
  }
],
"connectionServers":[
  {
    "fqdn":"cs1.horizon-1.local",
    "certificatePath":"/tmp/ConServCertificate8496572034532614986.p12",
    "certificatePassword":"leHOTEZ9",
    "state":"DEPLOYED",
    "deployDetails":{
      "ipAddress":"10.10.0.20",
      "gateway":"10.10.0.250",
      "subnetMask":"255.255.255.0",
      "vmName":"cs1",
      "computerName":"cs1"
    }
  }
],
"composerServers":[
  {
    "fqdn":"comp01.horizon-1.local",
    "composerServiceAccount":"horizon-1.local\\compsvc1",
    "composerServicePassword":"VMware123!",
    "sqlId":"sqlxt",
    "internalDbUsername":"internal-db-ext",
    "internalDbPassword":"idb-password",
    "dbName":"internal-db-ext",
    "certificatePath":"/tmp/CompServerCertificate6739599294661772272.p12",
    "certificatePassword":"IbZy8a0a",
    "vcFqdn":"vcenter-2.vrack.vsphere.local",
    "state":"DEPLOYED",
    "deployDetails":{
      "ipAddress":"10.10.0.22",
      "gateway":"10.10.0.250",
      "subnetMask":"255.255.255.0",
      "vmName":"comp01",
      "computerName":"comp01"
    }
  }
]

```

```

    ]
  }
],
"internalLbFqdn":"lb2.horizon-1.local",
"externalLbFqdn":"lb3.horizon-1.local",
"adminGroupName":"horizonadmins",
"uemDetails":{
  "uems":[
    {
      "fqdn":"uem.horizon-1.local",
      "configurationShare":"/share/drive",
      "profileArchiveShare":"profileshare",
      "state":"DEPLOYED",
      "deployDetails":{
        "dataDriveSizeGb":10,
        "configurationShareLocation":"/share/location",
        "profileArchiveShareLocation":"/profile/location",
        "ipAddress":"10.10.0.23",
        "gateway":"10.10.0.250",
        "subnetMask":"255.255.255.0",
        "vmName":"uem-vm",
        "computerName":"uem"
      }
    }
  ]
}
},
"appVolumesDetails":{
  "adminGroupName":"avadmins",
  "dbName":"appvol-db-ext",
  "sqlId":"sqlext",
  "lbFqdn":"lb1.horizon-1.local",
  "datastores":[
    {
      "id":"sddc-ds-1",
      "isPrimary":true
    }
  ]
},
"appVolumes":[
  {
    "fqdn":"appvol.horizon-1.local",
    "state":"DEPLOYED",
    "deployDetails":{
      "ipAddress":"10.10.0.24",
      "gateway":"10.10.0.250",
      "subnetMask":"255.255.255.0",
      "vmName":"app-vm",
      "computerName":"appvol"
    }
  }
]
}
}

```

View Horizon Domain Details

The Horizon Domain summary page displays all Horizon domains in your environment with a summary of the basic capacity information. You can select a domain to view the domain details, download the configuration, or expand or delete the domain.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domain**.
- 2 In the Horizon area, click **View Details**.

The summary page displays all Horizon domains and a summary of CPU, memory and storage usage.

- 3 Click the name of a Horizon domain.

The Horizon domain details page shows the status of the Horizon domain and the overall utilization summary for the domain. The Summary section displays the components of the Horizon domain (and their quantities).

- 4 Click the appropriate tab to see more information about service VMs, VI workload domains, and configuration details.

Tab	Information Displayed
Service VMs	Links to the administration consoles for each of the components that are used to manage the domain and the IP address of the component VMs.
VI	Summary and links to the VI workload domains associated with the Horizon domain.
Configuration Details	Configuration details for the Horizon domain.

Expand a Horizon Domain

You can expand a Horizon domain to add additional components (such as Connection Server), deploy optional components (such as App Volumes), or add VI workload domains to extend available desktop capacity.

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domain**.
- 2 In the Horizon area, click **View Details**.
- 3 Click  next to the Horizon domain whose configuration you want to export and click **Expand**.

The Horizon Domain Expansion window appears.

- 4 Proceed through the wizard pages and make the required updates. You can add additional elements, associate new VI workload domains, or add new components. For details on required information, see [Create a Horizon Domain](#).

Delete Horizon Domain

Deleting an active Horizon domain may affect desktops or RDSH server users, so be careful when deleting a Horizon domain.

A Horizon domain consists of two parts:

- VI desktop capacity in the VI workload domains
- Horizon management infrastructure

When you delete a Horizon domain, the Horizon management components (load balancers, Connection Servers, Composer Servers, and deployed optional components) are deleted. However, VI workload domains associated with the Horizon domain are not deleted, and desktops and RDSH servers in those VI workload domains are not deleted.

If you want to delete desktops, data, and management components, delete the VI workload domains associated with the Horizon domain before following this procedure. For information on deleting VI workload domains, see [Delete a VI Workload Domain](#).

Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domain**.
- 2 In the Horizon area, click **View Details**.
- 3 Click  next to the Horizon domain whose configuration you want to export and click **Delete**.
The Delete Domain window appears.
- 4 Type the name of the Horizon domain that you want to delete.
- 5 Click Delete Domain.

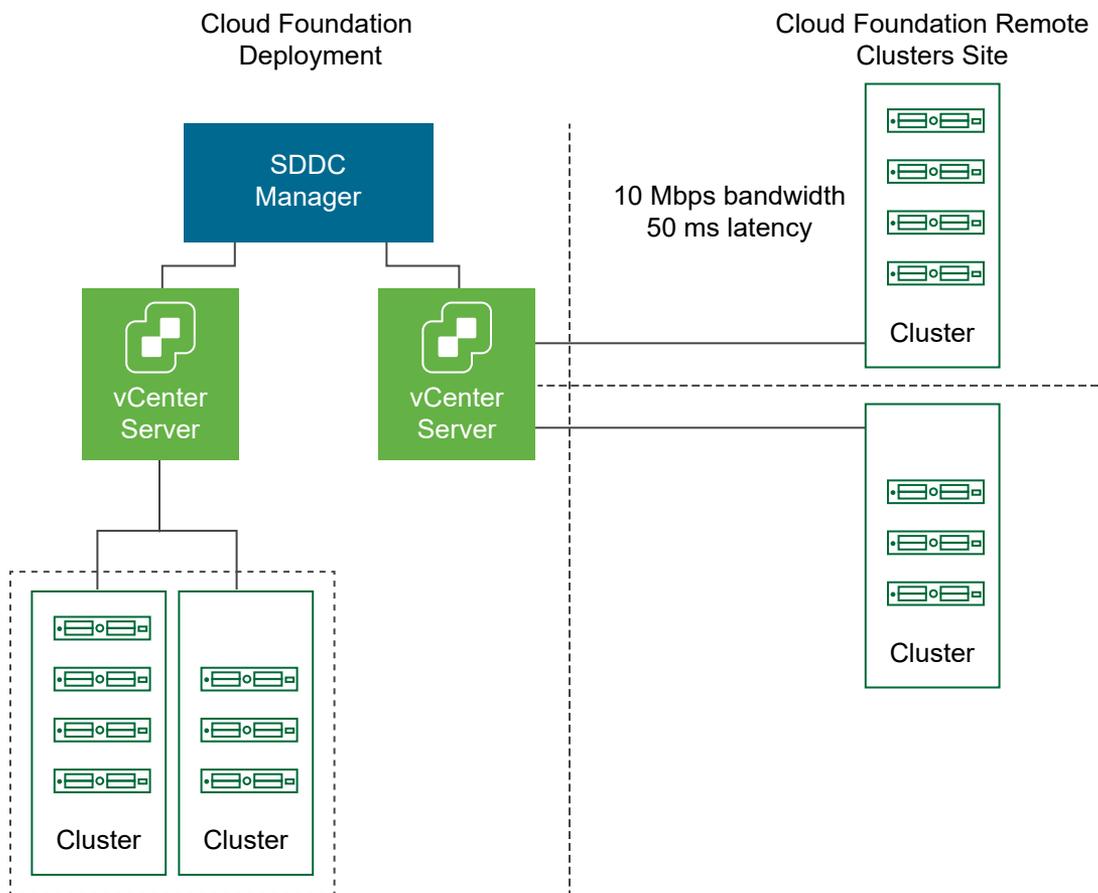
Results

If you deleted only the Horizon domain, the management components are deleted. Active Horizon sessions are interrupted and users will not be able to connect. The desktops and data are not deleted.

If you deleted the VI workload domains associated with the Horizon domain and the Horizon domain, all desktops, data, and management components are deleted.

Deploy a Workload Domain or Cluster at a Remote Location

With Cloud Foundation Remote Clusters, you can deploy a workload domain or cluster at a remote site through SDDC Manager. The remote workload domains and cluster are managed by the Cloud Foundation instance at the central site. You can perform a full-stack life cycle management for the remote sites from the central SDDC Manager.



This chapter includes the following topics:

- [Prerequisites for Cloud Foundation Remote Clusters](#)
- [Deploy a Workload Domain at a Remote Location](#)
- [Add a Cluster at a Remote Location](#)

Prerequisites for Cloud Foundation Remote Clusters

The remote site must meet the following prerequisites.

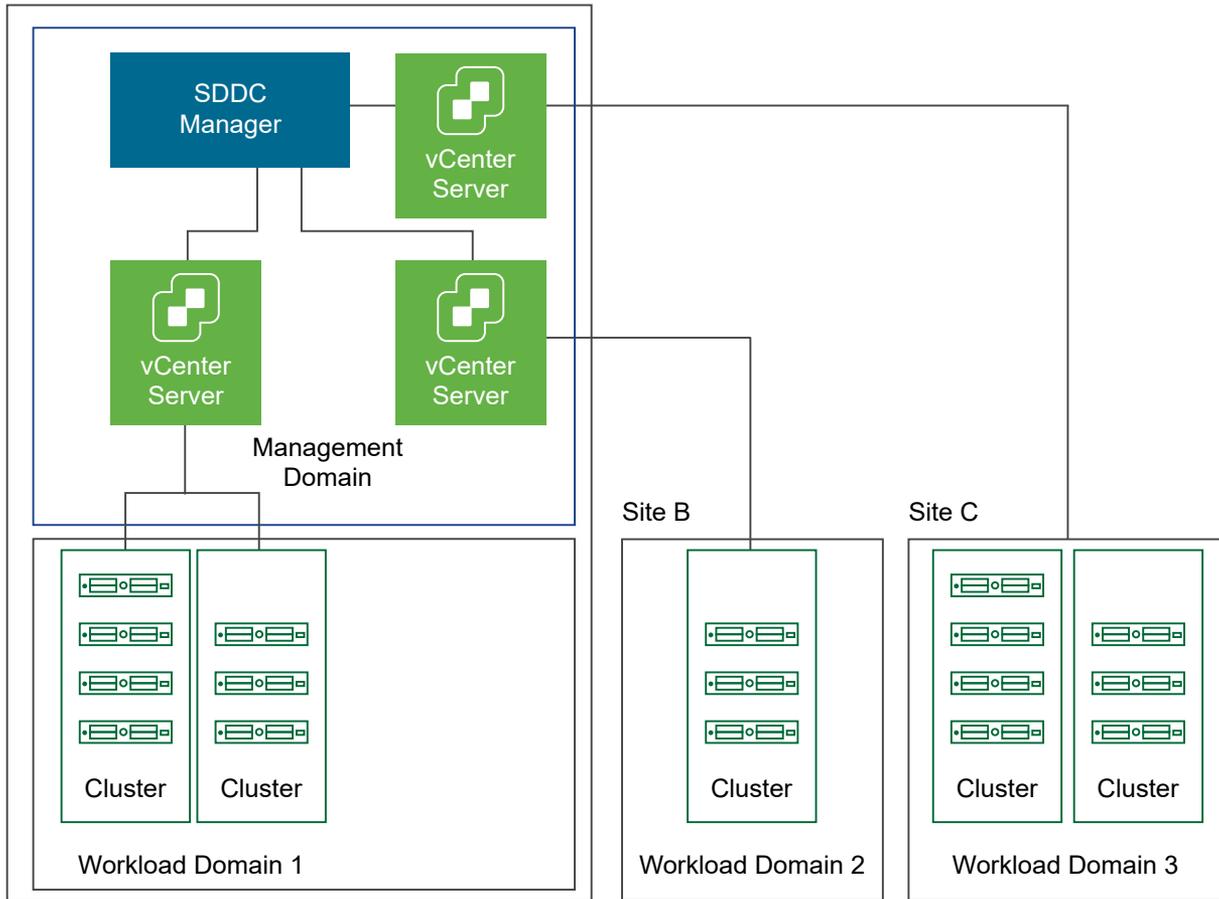
Prerequisites

- Dedicated WAN connectivity is required between central site and Cloud Foundation Remote Clusters site.
- Primary and secondary active WAN links are recommended for connectivity from the central site to the Cloud Foundation Remote Clusters site. The absence of WAN links can lead to a two-failure state, WAN link failure, or Edge node failure, which can result in unrecoverable VMs and application failure at the Cloud Foundation Remote Clusters site.
- Minimum bandwidth of 10 Mbps and latency of 50 Ms is required at the Cloud Foundation Remote Clusters site.
- The network at the Cloud Foundation Remote Clusters site must be able to reach the management network at the central site.
- DNS and NTP server must be available locally at or reachable from the Cloud Foundation Remote Clusters site
- DHCP Server must be available locally at the Cloud Foundation Remote Clusters site for NSX-T host overlay (host TEP) VLAN.

Deploy a Workload Domain at a Remote Location

When you deploy a workload domain at a remote location with Cloud Foundation Remote Clusters, you can perform all operations on the workload domain via the SDDC Manager at the central location.

Site A



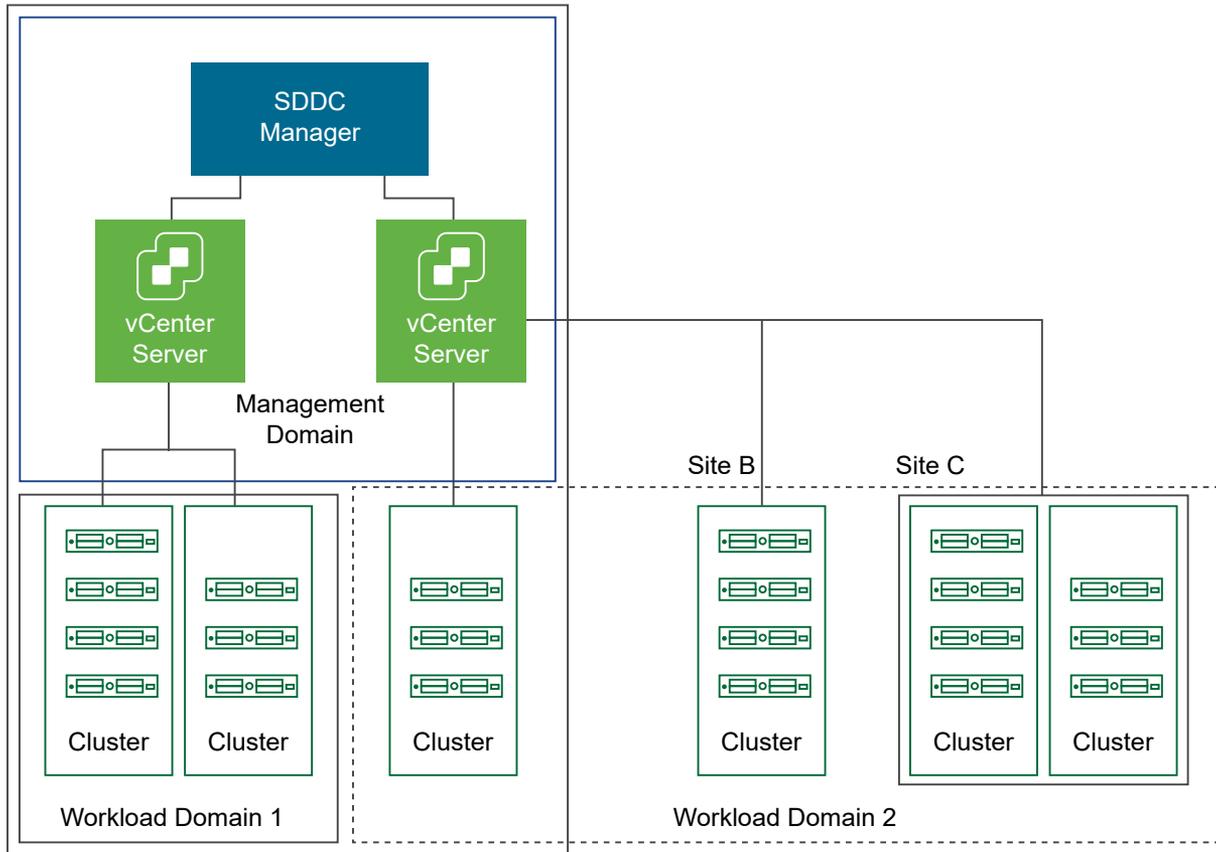
Before deploying the remote workload domain, you must ensure the Cloud Foundation Remote Clusters prerequisites are completed. See [Prerequisites for Cloud Foundation Remote Clusters](#) .

For information on deploying a workload domain, see [About VI Workload Domains](#).

Add a Cluster at a Remote Location

With Cloud Foundation Remote Clusters, you can expand a workload domain at the central location by adding a cluster at a remote location.

Site A



Before adding a cluster at a remote location, you must ensure the Cloud Foundation Remote Clusters prerequisites are completed. See [Prerequisites for Cloud Foundation Remote Clusters](#) . For information on adding a cluster, see [Add a Cluster to a Workload Domain](#).

vRealize Suite Products and Cloud Foundation

12

Using SDDC Manager, you can deploy vRealize Operations and vRealize Automation and connect them to workload domains in your Cloud Foundation system.

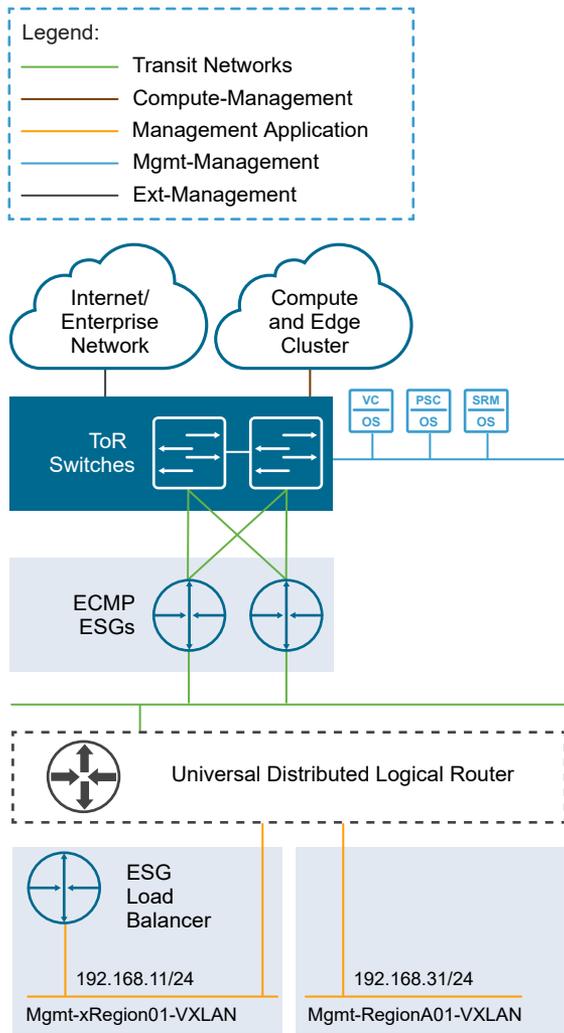
You can also enable vRealize Log Insight for all workload domains and expand the analytics cluster for vRealize Operations.

All vRealize Suite products require licenses purchases separately from Cloud Foundation.

Note For detailed information on prerequisites and preliminary procedures for adding vRealize Suite products to your Cloud Foundation deployment, see the *VMware Cloud Foundation Planning and Preparation Guide*.

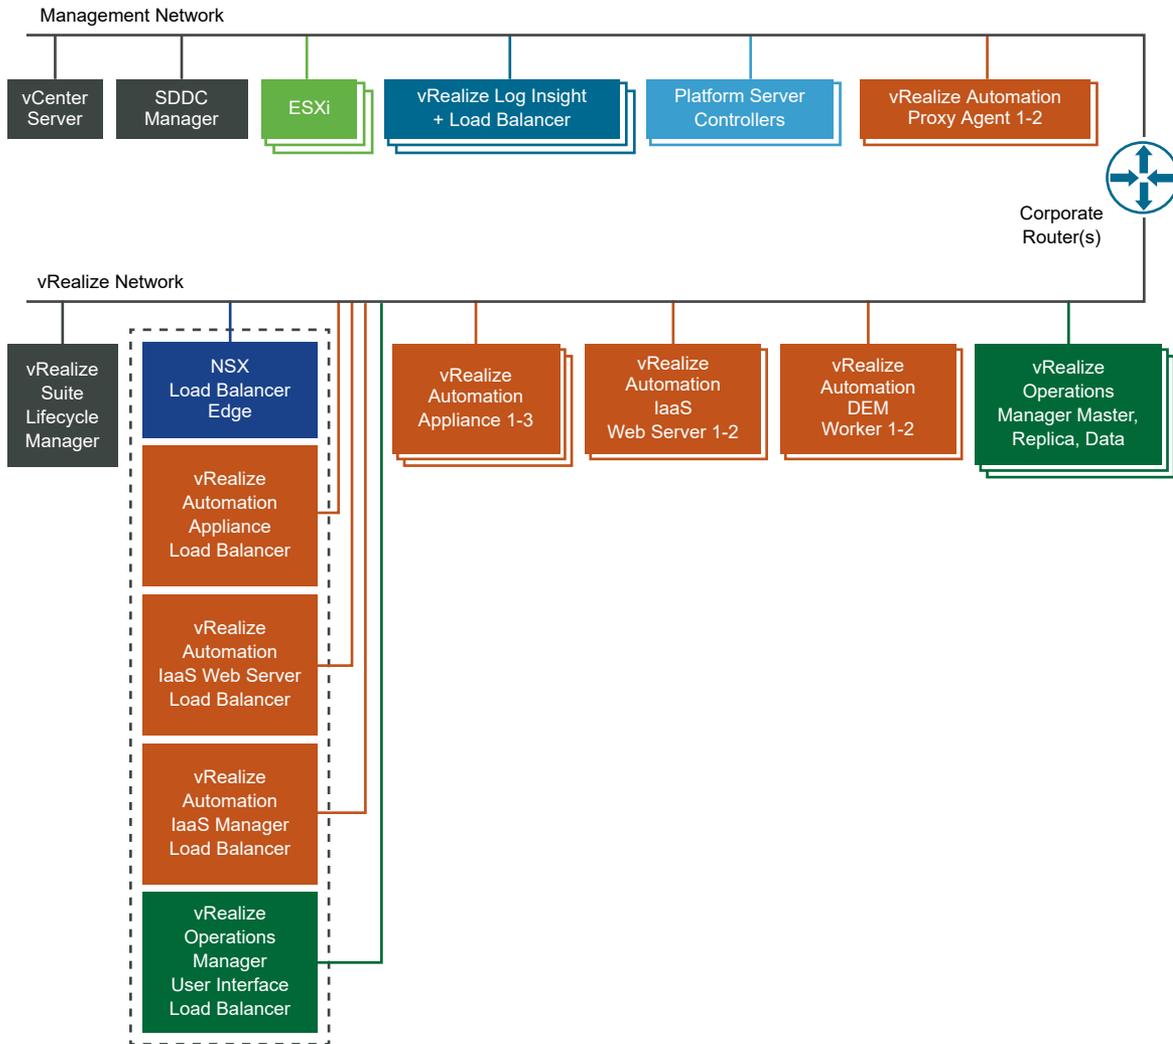
New installations of Cloud Foundation 3.9.1, deploy vRealize Suite components to virtual networks, called application virtual networks (AVNs) provided by VXLAN virtual wires. Cloud Foundation 3.9 deployed vRealize Suite components on a VLAN-backed distributed port group. After upgrading to 3.9.1, Cloud Foundation continues to use a VLAN-backed distributed port group for vRealize Suite components. To enable AVNs for an upgraded system and to migrate vRealize Suite components to the AVNs, contact VMware Support.

Figure 12-1. vRealize Suite Components with Application Virtual Networks



AVN	vRealize Component
Mgmt-RegionA01-VXLAN	vRealize Log Insight
Mgmt-RegionA01-VXLAN	vReze Automation Proxy Agents
Mgmt-xRegion01-VXLAN	vRealize Operations Manager
Mgmt-xRegion01-VXLAN	vRealize Automation
Mgmt-xRegion01-VXLAN	vRealize Suite Lifecycle Manager

Figure 12-2. vRealize Suite Components with a VLAN-backed distributed port group



Procedure

1 Deploy vRealize Suite Lifecycle Manager in Cloud Foundation

Before you can deploy vRealize Operations or vRealize Automation, you must deploy vRealize Suite Lifecycle Manager.

2 Adding vRealize Automation to Cloud Foundation

vRealize Automation provides a secure portal where authorized administrators, developers, and business users can request new IT services and manage specific cloud and IT resources according to business policies. Requests for IT service, including infrastructure, applications, desktops, and many others, are processed by using a common service catalog to provide a consistent user experience. SDDC Manager automates the deployment of vRealize Automation in Cloud Foundation.

3 Adding vRealize Operations to Cloud Foundation

vRealize Operations tracks and analyzes the operation of multiple data sources in Cloud Foundation by using specialized analytic algorithms. These algorithms help vRealize Operations learn and predict the behavior of every object it monitors. Users access this information by using views, reports, and dashboards. SDDC Manager automates the deployment of vRealize Operations in Cloud Foundation.

4 Connect vRealize Suite Products to Workload Domains in Cloud Foundation

You can connect your vRealize Automation and vRealize Operations deployments to workload domains in Cloud Foundation.

5 Add a Node to a vRealize Operations Analytics Cluster

When your vRealize Operations analytics cluster needs more resources, you can add additional nodes.

Deploy vRealize Suite Lifecycle Manager in Cloud Foundation

Before you can deploy vRealize Operations or vRealize Automation, you must deploy vRealize Suite Lifecycle Manager.

New VMware Cloud Foundation 3.9.1 installations use NSX Data Center for vSphere to create VXLAN-backed networks, called application virtual networks (AVNs), and deploy vRealize Suite components to these AVNs. Cloud Foundation 3.9 deployed vRealize Suite components on a VLAN-backed distributed port group. After upgrading to 3.9.1, Cloud Foundation continues to use VLAN-backed distributed port groups for vRealize Suite components.

If you want to migrate vRealize Suite components to AVNs, contact VMware Support.

Prerequisites

- Download the vRealize Suite Lifecycle Manager installation package from the VMware Depot to the local bundle repository. See [Chapter 13 Downloading an Install Bundle](#).
- Configure a new vRealize VLAN on the switches and verify that the vRealize subnet is routable to the management network. (Only required if you are using a VLAN-backed distributed port group for deploying vRealize Suite components.)
- Allocate an IP address for the vRealize Suite Lifecycle Manager virtual appliance and prepare forward/reverse DNS records.

Procedure

- 1 On the **SDDC Manager Dashboard**, navigate to **Administration > vRealize Suite**.

The **vRealize Suite** navigation appears, listing the vRealize Suite products available for your Cloud Foundation system deployment.

2 Click **vRealize Suite Lifecycle Manager**.

The **vRealize Suite Lifecycle Manager** page displays.

3 Click **Deploy**.

The **vRealize Lifecycle Manager Installation Prerequisites** page displays the prerequisites that you must complete before beginning the installation.

4 Review the readiness of each prerequisite and verify by selecting each adjacent check box.

When all the boxes are selected, the **Begin** button is activated.

5 Click **Begin**.

The **vRealize Lifecycle Manager Installation** wizard opens.

6 On the **Network Settings** page, enter or review the settings and click **Next** to continue.

If you are using AVNs, the **Network Settings** page displays the settings. You cannot edit these settings. If you are using a VLAN-backed distributed port group, enter the following settings:

Setting	Description
VLAN ID	Enter a valid VLAN ID between 0 and 4094 for the dedicated network.
Subnet Mask	Provide a valid subnet mask for the dedicated network.
Gateway	Provide a valid gateway address for the dedicated network.

7 On the **Virtual Appliance Settings** page, enter the settings and **Next** to continue.

Setting	Description
FQDN	Enter the FQDN for the vRealize Suite Lifecycle Manager virtual appliance.
System Administrator	Create and confirm a password for the vRealize Suite Lifecycle Manager system administrator (for example, admin@localuser). This is the credential that allows SDDC Manager to connect to the vRealize Suite Lifecycle Manager system.
SSH Root Account	Create and confirm a password for the vRealize Suite Lifecycle Manager virtual appliance root account.

8 On the **Review Summary** page, review the installation configuration settings.**9** Click **Finish**.

SDDC Manager validates the inputs and reports any errors or warnings.

Note If necessary, you can use the **Back** button to return to preceding pages and modify settings.

- 10 Address any validation issues and then click **Finish**.

The **vRealize Suite Lifecycle Manager** page displays with the following message: **Deployment in progress**. If the deployment fails, this page displays a deployment status of **Failed**. In this case, you can **Restart Task** or **Uninstall**.

- 11 (Optional) Click **View Status in Tasks** to view the details of the deployment in progress or a deployment failure.

The **Tasks** panel opens at the bottom page. You can open individual tasks to view details.

- 12 (Optional) After the successful deployment of vRealize Suite Lifecycle Manager, click the vRealize Suite Lifecycle Manager link below the page title.

The vRealize Suite Lifecycle Manager user interface opens in a new browser tab.

What to do next

You can now deploy vRealize Operations or vRealize Automation.

Adding vRealize Automation to Cloud Foundation

vRealize Automation provides a secure portal where authorized administrators, developers, and business users can request new IT services and manage specific cloud and IT resources according to business policies. Requests for IT service, including infrastructure, applications, desktops, and many others, are processed by using a common service catalog to provide a consistent user experience. SDDC Manager automates the deployment of vRealize Automation in Cloud Foundation.

Add a vRealize Automation License Key to Cloud Foundation

Before you can deploy vRealize Automation for use with Cloud Foundation you must add a license.

Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Click **+ License Key**.
- 3 Select **VMware vRealize Automation** as the product name.
- 4 Type the license key.

You can enter a license key for vRealize Automation or a suite license for a product that includes vRealize Automation. For example, vCloud Suite or vRealize Suite.

- 5 Type a description for the license.

If you have multiple license keys for a product, the description can help in identifying the license.

- 6 Click **Add**.

Deploy vRealize Automation in Cloud Foundation

You deploy vRealize Automation in Cloud Foundation using the SDDC Manager user interface.

New VMware Cloud Foundation 3.9.1 installations use NSX Data Center for vSphere to create VXLAN-backed networks, called application virtual networks (AVNs), and deploy vRealize Suite components to these AVNs. Cloud Foundation 3.9 deployed vRealize Suite components on a VLAN-backed distributed port group. After upgrading to 3.9.1, Cloud Foundation continues to use VLAN-backed distributed port groups for vRealize Suite components.

If you want to migrate vRealize Suite components to AVNs, contact VMware Support.

Prerequisites

- Deploy vRealize Suite Lifecycle Manager. See [Deploy vRealize Suite Lifecycle Manager in Cloud Foundation](#).
- Verify you have a valid license key for vRealize Automation, which is purchased separately from Cloud Foundation. You must add the license key to Cloud Foundation before deploying vRealize Automation. See [Add a vRealize Automation License Key to Cloud Foundation](#).
- Verify you have downloaded the vRealize Suite bundles from the VMware Depot. See [Chapter 13 Downloading an Install Bundle](#).
- Verify that IP allocation and forward/reverse DNS records are prepared for the vRealize Automation components.
- Verify you have created the required Active Directory (AD) service account for vRealize Automation.
- Verify that you have configured a certificate authority in SDDC Manager. See [Configure a Microsoft Certificate Authority](#).
- Verify the multi-SAN certificate and private key generated by a trusted certificate authority is available for vRealize Automation.
- Verify Microsoft SQL Server is properly deployed and configured for vRealize Automation.
- Verify you have created and exported a Microsoft Windows Server OVA template for the vRealize Automation IaaS components.

For more information, see the *VMware Cloud Foundation Planning and Preparation Guide*.

Procedure

- 1 On the **SDDC Manager Dashboard**, navigate to **Administration > vRealize Suite**.

The **vRealize Suite** navigation appears, listing the vRealize Suite products available for your Cloud Foundation system deployment.

- 2 Click **vRealize Automation**.

The **vRealize Automation** page displays.

3 Click **Deploy**.

The **vRealize Automation Installation Prerequisites** page displays the prerequisites that you must complete before beginning the installation.

4 Review the readiness of each prerequisite and verify by selecting each adjacent check box.

When all the boxes are selected, the **Begin** button is activated.

5 Click **Begin**.

The **vRealize Automation Installation** wizard opens.

6 On the **Deployment Details** page, enter the settings and **Next** to continue.

All settings are required.

Setting	Description
vRealize Automation License Key	<p>License Key</p> <p>Select a valid license for vRealize Automation. This license may be for vRealize Automation, vRealize Suite, or vCloud Suite. If no key is available, you can add one in Administration > Licensing.</p>
Certificate Details	<p>Certificate Chain</p> <p>Enter the full certificate chain, including each -----BEGIN CERTIFICATE----- header and -----END CERTIFICATE----- footer.</p> <p>The certificate chain is a combination of the server certificate and the root CA certificate, in that order.</p> <p>Certificate Private Key</p> <p>Enter the private key for the certificate, including the -----BEGIN RSA PRIVATE KEY----- header and the -----END RSA PRIVATE KEY----- footer.</p> <p>Create/Confirm Passphrase</p> <p>Create and confirm a passphrase for the certificate protection on the vRealize Automation IaaS Windows servers.</p>
IaaS Windows Template	<p>Select one of the following options from the drop-down options:</p> <p>Upload OVA Template</p> <p>Select to upload a new Windows Server OVA template. Click Upload to navigate to and upload the OVA template. The filename must be less than 60 characters in length, including the extension.</p> <p>Use Existing OVA Template</p> <p>Select to use and specify an existing Windows Server OVA template. The OVA template path automatically displays the path to an existing template uploaded to SDDC Manager using SCP, if any.</p> <p>The file path must have the following permissions:</p> <p>owner: vcf_commonsvcs</p> <p>group: vcf</p> <p>The directories in the path must be readable and executable for the user and the group.</p> <p>For example:</p> <pre>chmod 0750 -R /upload chown vcf_commonsvcs:vcf -R /upload</pre>

- 7** If you are using AVNs, the **Network Settings** page displays read-only information about the application virtual networks that vRealize Automation will use. Click **Next**.

8 On the **FQDNs** page, enter the settings and **Next** to continue.

Important The installation derives the Active Directory domain name for the computer account from the DNS suffix provided in the FQDN for each vRealize Automation IaaS component. For example, an FQDN of **vra01iws01a.rainpole.local** derives the Active Directory domain **rainpole.local**. If the DNS suffix is different from Active Directory domain name, the installation will be unsuccessful. For more information, see [Knowledge Base article 59128](#).

Note FQDNs must start with a letter and can contain only letters, digits, dashes, and underscores.

Setting	Description
vRealize Automation Appliances	<p>Appliance 1 Enter the FQDN as provided in the certificate.</p> <p>Appliance 2 Enter the FQDN as provided in the certificate.</p> <p>Appliance 3 Enter the FQDN as provided in the certificate.</p>
IaaS Web Servers	<p>IaaS Web Server 1 Enter the FQDN as provided in the certificate.</p> <p>IaaS Web Server 2 Enter the FQDN as provided in the certificate.</p> <p>Note Host names for Windows IaaS VMs should be 15 characters or less due to limitations in the Windows OS. If the host names are longer they will be trimmed during the installation and installation will fail.</p>
IaaS Manager Service and DEM Orchestrators	<p>IaaS Manager 1 Enter the FQDN as provided in the certificate.</p> <p>IaaS Manager 2 Enter the FQDN as provided in the certificate.</p> <p>Note Host names for Windows IaaS VMs should be 15 characters or less due to limitations in the Windows OS. If the host names are longer they will be trimmed during the installation and installation will fail.</p>

Setting	Description
DEM Workers	<p>DEM Worker 1</p> <p>Enter the FQDN as provided in the certificate.</p> <p>DEM Worker 2</p> <p>Enter the FQDN as provided in the certificate.</p> <p>Note Host names for Windows IaaS VMs should be 15 characters or less due to limitations in the Windows OS. If the host names are longer they will be trimmed during the installation and installation will fail.</p>
Proxy Agents	<p>Proxy Agent 1</p> <p>Enter the FQDN.</p> <p>Proxy Agent 2</p> <p>Enter the FQDN.</p> <p>Note Host names for Windows IaaS VMs should be 15 characters or less due to limitations in the Windows OS. If the host names are longer they will be trimmed during the installation and installation will fail.</p>
vRealize Suite Lifecycle Manager	<p>Hostname</p> <p>Displays the FQDN.</p>
Load Balancers	<p>When you deploy vRealize Automation, an NSX Edge Service Gateway is deployed as a one-armed load balancer. This load balancer is shared between vRealize Operations and vRealize Automation. If you already deployed vRealize Operations, the NSX Edge Service Gateway FQDN displays as read-only.</p> <p>NSX Edge Services Gateway</p> <p>Displays the FQDN.</p> <p>IaaS Web Server Virtual Server</p> <p>Enter the FQDN as provided in the certificate.</p> <p>IaaS Manager Virtual Server</p> <p>Enter the FQDN as provided in the certificate.</p> <p>vRealize Automation Appliance Virtual Server</p> <p>Enter the FQDN as provided in the certificate.</p>
Microsoft SQL Server	<p>Hostname</p> <p>Provide the FQDN for the Microsoft SQL Server virtual appliance.</p>

9 On the **Account Information** page, enter the settings and **Next** to continue.

Note All settings are required.

Setting	Description
Active Directory	<p>Use these settings to provide the service account that is used for services on the IaaS VMs. This account must have administrative permissions to join Windows VMs to Active Directory.</p> <p>Username</p> <p>Provide the service account user name in the "domain\username" format.</p> <p>Password / Confirm Password</p> <p>Provide and confirm a valid password.</p>
Microsoft SQL Server	<p>Use these settings to create the connection to the database.</p> <p>Database Name</p> <p>Specify the case-sensitive database name.</p> <p>Username</p> <p>Specify the database owner user name. This setting is optional. If no user name is specified, Cloud Foundation applies the Active Directory account used when the database was joined to Active Directory.</p> <p>Password / Confirm Password</p> <p>Provide and confirm a valid password for the specified user. This is required only if you also provide a username, as described above.</p>
Local Tenant Administrator	<p>Use these settings to create a new user for the default vRealize Automation tenant. This user will be assigned the Tenant Administrator role.</p> <p>First Name</p> <p>Enter the administrator's first name.</p> <p>Last Name</p> <p>Enter the administrator's last name.</p> <p>Email</p> <p>Enter the administrator's email.</p> <p>Username</p> <p>Define a user name for the tenant administrator.</p> <p>Create Password / Confirm Password</p> <p>Create and confirm a password for the tenant administrator.</p>
Windows Template Local Administrator	<p>Create Password / Confirm Password</p> <p>Create and confirm the local administrator password for the Windows system that is deployed through the Windows IaaS VM template.</p>

Setting	Description
Default Tenant Administrator	<p>Create Password / Confirm Password</p> <p>Create and confirm the password for the vRealize Automation system administrator. This is the credential that allows SDDC Manager to connect to the vRealize Automation system.</p>
vRealize Automation SSH Root Account	<p>Create Password / Confirm Password</p> <p>Create and confirm a password for the vRealize Automation virtual appliance root account.</p>

- 10** On the **Review Summary** page, review a summary of the installation configuration settings. This page displays any validation errors that require attention.

Note If necessary, you can use the **Back** button to return to preceding pages and modify settings. You can also proceed without validation.

- 11** Click **Finish**.

The **vRealize Automation** page displays with the following message: **Deployment in progress**. If the deployment fails, this page displays a deployment status of **Failed**. In this case, you can **Retry** or **Uninstall**.

Important The uninstall operation does not remove the computer accounts from Active Directory. As a result, this could cause a reinstallation to fail. Manually remove the computer accounts from Active Directory and recreate the Microsoft SQL Server database for vRealize Automation. See the *VMware Cloud Foundation Planning and Preparation Guide*.

- 12** (Optional) Click **View Status in Tasks** to view the details of the deployment in progress or a deployment failure.

The **Tasks** panel opens at the bottom page. You can open individual tasks to view details.

- 13** (Optional) After the successful deployment of vRealize Automation, click the vRealize Automation link below the page title.

The vRealize Automation user interface opens in a new browser tab.

Results

After the successful deployment of vRealize Automation, the **vRealize Automation** page in **SDDC Manager > Administration > vRealize Suite** displays an **ACTIVE** status and displays controls that enable you to connect vRealize Automation to workload domains.

What to do next

You must manually start the vRealize Orchestrator configuration service. See [Start the vRealize Orchestrator Configurator Service in Cloud Foundation](#).

Post-Deployment Tasks for vRealize Automation in Cloud Foundation

After you complete the procedure to add vRealize Automation to Cloud Foundation, verify that the following configurations are established.

Start the vRealize Orchestrator Configurator Service in Cloud Foundation

After deploying vRealize Automation in Cloud Foundation, you must manually start the vRealize Orchestrator Configurator service to access the vRealize Orchestrator configuration interface.

Procedure

- 1 Log in to the first vRealize Automation appliance by using Secure Shell (SSH) client to configure the embedded vRealize Orchestrator Configurator service.
- 2 Verify that the vRealize Orchestrator user interface service is running.
 - a Run the following command to verify that the service is set to automatically start.

```
chkconfig vco-configurator
```

- b If the service reports Off, run the following command to enable an automatic restart of the vRealize Orchestrator Configurator service upon subsequent reboots of the vRealize Automation appliance.

```
chkconfig vco-configurator on
```

- c Verify the status of the vRealize Orchestrator Configurator service by running the following command .

```
service vco-configurator status
```

- d Repeat the procedure to configure vRealize Orchestrator for the other vRealize Automation appliances.

Add Tenant Administrator Permissions for vRealize Orchestrator

After deploying vRealize Automation, you must add the tenant administrator role to the vRealize Orchestrator **vcoadmins** group.

vRealize Orchestrator allows administrators and architects to develop complex automation tasks by using the workflow designer, and to access and run the workflows from vRealize Automation.

Procedure

- 1 Log in to the vRealize Automation console (<https://vra-appliance-FQDN>) as administrator using the password you specified during deployment.
- 2 Click **Tenants**.
- 3 Select `vsphere.local` and click **Edit**.

- 4 Click **Administrators**.
- 5 Enter **vcoadmins** in the Tenant administrators search box.
- 6 Select vcoadmins (vcoadmins@vsphere.local) to grant the tenant administrator role to that user group.
- 7 Click **Finish**.
- 8 Log in to the first vRealize Automation appliance using Secure Shell (SSH).
- 9 Run the following commands to restart the vRealize Orchestrator services.

```
service vco-server restart
```

```
service vco-configurator restart
```

- 10 Repeat steps 8 and 9 for the remaining vRealize Automation appliances.

Create VM Groups to Define the Startup Order of vRealize Automation in Cloud Foundation

Define the startup order of vRealize Automation components with VM Groups. The startup order ensures that vSphere HA powers on virtual machines in the correct order,

Procedure

- 1 On the SDDC Manager Dashboard, select **Inventory > Workload Domains** from the **Navigation** pane.
- 2 Click on the **MGMT** Management Domain.
- 3 Select the **Services** tab on the Management Domain.
- 4 Under the **VMware Cloud Foundation Components** section, click on the link for the vCenter Server.

A new browser window launch the landing page for the vSphere Web Client.

- 5 On the **Welcome to VMware vSphere** browser windows, click the link for **vSphere Web Client (Flash)**.

The vSphere Web Client will open.

- 6 In the **Navigator**, select **Host and Clusters** and expand the tree for the Management Domain vCenter Server instance..
- 7 Create a VM Group for the vRealize Automation IaaS database.
 - a Select the Management Domain cluster and click the **Configure** tab.
 - b On the **Configure** page, click **VM/Host Groups**.
 - c On the **VM/Host Groups** page, click the **Add** button.

- d In the **Create VM/Host Group** dialog box, enter **vRealize Automation IaaS Database** in the **Name** text box, select **VM Group** from the **Type** drop-down menu, and click the **Add** button.
- e In the **Add VM/Host Group Member** dialog box, select virtual machine for the Microsoft SQL Server (for example, **vra01mssql01**) and click **OK**.
- f Click **OK** to save the VM/Host Group.

8 Repeat *Step 3* to create the following VM/Host Groups

VM/Host Group Name	VM/Host Group Member
vRealize Automation Virtual Appliances	vRealize Automation Appliance 1
-	vRealize Automation Appliance 2
-	vRealize Automation Appliance 3
vRealize Automation IaaS Web Servers	vRealize Automation IaaS Web Server 1
-	vRealize Automation IaaS Web Server 2
vRealize Automation IaaS Manager Servers	vRealize Automation IaaS Manager Server 1
-	vRealize Automation IaaS Manager Server 2
vRealize Automation IaaS DEM Workers	vRealize Automation IaaS DEM Worker 1
-	vRealize Automation IaaS DEM Worker 2
vRealize Automation Proxy Agents	vRealize Automation Proxy Agent 1
-	vRealize Automation Proxy Agent 2

- 9** Create a rule to power on the vRealize Automation IaaS database virtual machine before the vRealize Automation virtual appliances and vRealize Automation IaaS virtual machines..
 - a Select the Management Domain cluster and click the **Configure** tab.
 - b On the **Configure** page, click **VM/Host Rules**.
 - c On the **VM/Host Rules** page, click the **Add** button.
 - d In the **Create VM/Host Rule** dialog box, enter **SDDC Cloud Management Platform 01** in the **Name** text box, ensure the **Enable Rule** check box is selected, and select **Virtual Machines to Virtual Machines** from the **Type** drop-down menu.
 - e Select **vRealize Automation IaaS Database** from the **First restart VMs in VM group** drop-down menu.
 - f Select **vRealize Automation Virtual Appliances** from the **Then restart VMs in VM group** drop-down menu, and click **OK**.

- 10** Repeat *Step 5* to create the following VM/Host Rules to ensure the correct restart order for your Cloud Management Platform.

VM/Host Rule Name	First restart VMs in VM group	Then restart VMs in VM group
SDDC Cloud Management Platform 02	vRealize Automation Virtual Appliances	vRealize Automation IaaS Web Servers
SDDC Cloud Management Platform 03	vRealize Automation IaaS Web Servers	vRealize Automation IaaS Manager Servers
SDDC Cloud Management Platform 04	vRealize Automation IaaS Manager Servers	vRealize Automation IaaS DEM Workers
SDDC Cloud Management Platform 05	vRealize Automation IaaS Manager Servers	vRealize Automation Proxy Agents

Results

Adding vRealize Operations to Cloud Foundation

vRealize Operations tracks and analyzes the operation of multiple data sources in Cloud Foundation by using specialized analytic algorithms. These algorithms help vRealize Operations learn and predict the behavior of every object it monitors. Users access this information by using views, reports, and dashboards. SDDC Manager automates the deployment of vRealize Operations in Cloud Foundation.

Add a vRealize Operations License Key to Cloud Foundation

Before you can deploy vRealize Operations for use with Cloud Foundation you must add a license.

Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Click **+ License Key**.
- 3 Select **VMware vRealize Operations** as the product name.

- 4 Type the license key.

You can enter a license key for vRealize Operations or a suite license for a product that includes vRealize Operations. For example, vCloud Suite or vRealize Suite.

- 5 Type a description for the license.

If you have multiple license keys for a product, the description can help in identifying the license.

- 6 Click **Add**.

Deploy vRealize Operations in Cloud Foundation

You deploy vRealize Operations in Cloud Foundation using the SDDC Manager user interface.

New VMware Cloud Foundation 3.9.1 installations use NSX Data Center for vSphere to create VXLAN-backed networks, called application virtual networks (AVNs), and deploy vRealize Suite components to these AVNs. Cloud Foundation 3.9 deployed vRealize Suite components on a VLAN-backed distributed port group. After upgrading to 3.9.1, Cloud Foundation continues to use VLAN-backed distributed port groups for vRealize Suite components.

If you want to migrate vRealize Suite components to AVNs, contact VMware Support.

Prerequisites

- Deploy vRealize Suite Lifecycle Manager. See [Deploy vRealize Suite Lifecycle Manager in Cloud Foundation](#).
- Verify you have a valid license key for vRealize Operations, which is purchased separately from Cloud Foundation. You must add the license key to Cloud Foundation before deploying vRealize Operations. See [Add a vRealize Operations License Key to Cloud Foundation](#).
- Verify you have downloaded the vRealize Suite bundles from the VMware Depot. The bundle is obtained separately from the Cloud Foundation installation download. See [Chapter 13 Downloading an Install Bundle](#).

- Verify that IP allocation and forward/reverse DNS records are prepared the vRealize Operations components.

For more information, see the *VMware Cloud Foundation Planning and Preparation Guide*.

- Verify that you have determined the size of the vRealize Operations deployment to provide enough resources to accommodate the analytics operations for monitoring the expected number of workloads and SDDC management packs in the Cloud Foundation system.

For more information, use the online [vRealize Operations Sizing](#) utility.

Procedure

- 1 On the **SDDC Manager Dashboard**, navigate to **Administration > vRealize Suite**.

The **vRealize Suite** navigation appears, listing the vRealize Suite products available for your Cloud Foundation system deployment.

- 2 Click **vRealize Operations**.

The **vRealize Operations** page displays.

- 3 Click **Deploy**.

The **vRealize Operations Installation Prerequisites** page displays the prerequisites that you must complete before beginning the installation.

- 4 Review the readiness of each prerequisite and verify by selecting each adjacent check box.

When all the boxes are selected, the **Begin** button is activated.

- 5 Click **Begin**.
- 6 On the **Deployment Details** page, enter the settings and **Next** to continue.

All settings are required.

Setting	Description
License Key	Select a valid license for vRealize Operations. This license may be for vRealize Operations, vRealize Suite, or vCloud Suite. If no key is available, you can add one in Administration > Licensing .
High Availability	Optionally, move the button to green to deploy vRealize Operations with high availability configured.
Node Size	Select a Node Size based on your requirements. Note If you enable High Availability , you must specify a Node Size of Medium or larger.
Node Count	Select a Node Count for based on your requirements. Note If you enable High Availability , you must specify a Node Count of 2 or more.

Note The **Node Size** limits the number of nodes you can specify. Review the vRealize Operations Sizing Guidelines in VMware Knowledge Base article [54370](#).

- 7 If you are using AVNs, the **Network Settings** page displays read-only information about the application virtual networks that vRealize Operations will use. Click **Next**.
- 8 On the **FQDNs** page, enter the settings and **Next** to continue.

Note FQDNs must start with a letter and can contain only letters, digits, dashes, and underscores.

Setting	Description
Load Balancers	<p>When you deploy vRealize Operations, an NSX Edge Service Gateway is deployed as a one-armed load balancer. This load balancer is shared between vRealize Operations and vRealize Automation. If you already deployed vRealize Automation, the NSX Edge Service Gateway FQDN displays as read-only.</p> <p>NSX Edge Service Gateway</p> <p>Enter the FQDN.</p> <p>vRealize Operations</p> <p>Enter the FQDN for the vRealize Operations virtual server on the NSX Edge load balancer.</p>
vRealize Operations Nodes	<p>Node 1</p> <p>Enter the FQDN.</p> <p>Node n</p> <p>Enter the FQDN for each Node n.</p> <p>For example, if you specify 3 nodes:</p> <ul style="list-style-type: none"> ■ Without High Availability enabled. <ul style="list-style-type: none"> ■ Node 1 = Master Node FQDN ■ Node 2 = Data Node 1 FQDN ■ Node 3 = Data Node 2 FQDN ■ With High Availability enabled. <ul style="list-style-type: none"> ■ Node 1 = Master Node FQDN ■ Node 2 = Replica Node FQDN ■ Node 3 = Data Node FQDN
vRealize Suite Lifecycle Manager	<p>vRealize Suite Lifecycle Manager</p> <p>Displays the FQDN.</p>

- 9 On the **Account Information** page, enter the settings and **Next** to continue.

Note All settings are required.

Setting	Description
vRealize Operations Systems Administrator	<p>Create Password / Confirm Password</p> <p>Create and confirm a password for the vRealize Operations system administrator. This is the credential that allows SDDC Manager to connect to the vRealize Operations system.</p>

- 10 On the **Review Summary** page, review a summary of the installation configuration settings. This page displays any validation errors that require attention.

Note If necessary, you can use the **Back** button to return to preceding pages and modify settings. You can also proceed without validation.

11 Click **Finish**.

The **vRealize Operations** page displays with the following message: Deployment in progress.

If the deployment fails, this page displays a deployment status of **Failed** and prompts you to **Uninstall**.

Click **Uninstall** to return to the **vRealize Operations** page. Confirm your configuration settings, and retry the deployment operation.

12 After deploying vRealize Operations on Cloud Foundation, you must replace the security certificate.

See [Chapter 4 Certificate Management](#).

13 (Optional) Click **View Status in Tasks** to view the details of the deployment in progress or a deployment failure.

The **Tasks** panel opens at the bottom page. You can open individual tasks to view details.

14 (Optional) After the successful deployment of vRealize Operations, click the vRealize Operations link below the page title.

The vRealize Operations user interface opens in a new browser tab.

Results

After the successful deployment of vRealize Operations, the **vRealize Operations** page in **SDDC Manager > Administration > vRealize Suite** displays an **ACTIVE** status and displays controls that enable you to connect vRealize Operations to workload domains.

Post-Deployment Tasks for vRealize Operations in Cloud Foundation

After you complete the procedure to add vRealize Operations to Cloud Foundation, verify that the following configurations are established.

Important In addition to the below procedures, after deploying vRealize Operations on Cloud Foundation, you must replace the security certificate. See [Chapter 4 Certificate Management](#).

Configure SSL Passthrough for vRealize Operations Manager

By default, the vRealize Operations Manager node's load balancer is configured for SSL Termination. If you plan to use a custom certificate with vRealize Operations Manager, it is recommended that you replace the certificate on the vRealize Operations Manager cluster and configure the load balancer for SSL Passthrough.

Prerequisites

Verify that you have successfully replaced the vRealize Operations Manager certificate using the workflow described in [Chapter 4 Certificate Management](#).

Procedure

- 1 Log in into the management vCenter Server and navigate to **Home > Networking & Security**.
- 2 Select **NSX Edges** in the Navigator.
- 3 Confirm that the IP address in the **NSX Manager** field is identical to the IP address for the NSX Manager for the management domain in Cloud Foundation.
- 4 Double-click the NSX Edge labeled **vrealize-edge**.
- 5 Select the **Manage** tab, then the **Load Balancer** tab.
- 6 Open **Application Profiles**.
- 7 Find and click the profile named **vrops-https**, and click **Edit**.
- 8 In the **Application Profile Type** drop-down menu, select **SSL Passthrough** and click **OK**.
- 9 Log into the vRealize Operations Manager Master node as root via SSH or Console.
- 10 Open `/usr/lib/vmware-vcopsuite/utilities/conf/vcops-apache.conf` in a text editor.
- 11 Find the `ServerName ${VCOPS_APACHE_SERVER_NAME}` line and insert a new line after it.
- 12 On the new line enter the following:

```
ServerAlias vrops-lb.vrack.vsphere.local vrops-master.vrack.vsphere.local
```

Replace `vrops-lb.vrack.vsphere.local` with the FQDN of vRealize Operations Manager load balancer and replace `vrops-master.vrack.vsphere.local` with the FQDN of the vRealize Operations Manager master node.

- 13 Save and close the file.
- 14 Restart the `apache2` service:

```
service apache2 restart
```

- 15 Repeat steps 9-14 for all nodes in the vRealize Operations Manager cluster.

Connect vRealize Suite Products to Workload Domains in Cloud Foundation

You can connect your vRealize Automation and vRealize Operations deployments to workload domains in Cloud Foundation.

When connected, vRealize Automation and vRealize Operations monitor and collect data on the workload domains in Cloud Foundation.

Note This version of Cloud Foundation does not support connecting vRealize Automation to NSX-T workload domains through the SDDC Manager Dashboard. Use the vRealize Automation console instead.

By default, the management workload domain is connected to vRealize Operations. You can also enable log collection by enabling vRealize Log Insight within SDDC Manager.

Note You can create only one connection at a time.

Important Once you enable a connection between vRealize Automation and a workload domain, and then complete the connection wizard, you cannot disable the connection.

Prerequisites

- Verify that one or more workload domains has been created.
- Verify that vRealize Automation and vRealize Operations are deployed and operational.

Procedure

1 [Connect vRealize Suite Products to Workload Domains in Cloud Foundation](#)

You can connect vRealize Operations and vRealize Automation product deployments in Cloud Foundation to your workload domains in the SDDC Manager user interface

2 [Connect Workload Domains to vRealize Suite Products in Cloud Foundation](#)

You can connect workload domains to vRealize Operations and vRealize Automation product deployments in Cloud Foundation using the SDDC Manager user interface.

3 [Enable vRealize Log Insight in Cloud Foundation](#)

You can connect vRealize Log Insight in Cloud Foundation to all workload domains in the SDDC Manager user interface.

Connect vRealize Suite Products to Workload Domains in Cloud Foundation

You can connect vRealize Operations and vRealize Automation product deployments in Cloud Foundation to your workload domains in the SDDC Manager user interface

Note This version of Cloud Foundation does not support connecting vRealize Automation to NSX-T workload domains through the SDDC Manager Dashboard. Use the vRealize Automation console instead.

Prerequisites

- Before you can connect the management domain or workload domains to vRealize Operations, it must be deployed. For more information, see [Deploy vRealize Operations in Cloud Foundation](#).
- Before you can connect workload domains to vRealize Automation, it must be deployed. For more information, see [Deploy vRealize Automation in Cloud Foundation](#).

Procedure

- 1 On the **SDDC Manager Dashboard**, navigate to **Administration > vRealize Suite**.

2 To connect your vRealize Operations deployment to workload domains:

- a Select **vRealize Operations**.
- b Under **Connect Workload Domains...**, click **Connect**.

The **Connect to Workload Domains** wizard opens to the **Modify Connection** page. This page lists all currently configured workload domains and enables you to connect vRealize Operations to each

3 To connect your vRealize Automation deployment to workload domains:

- a Select **vRealize Automation**.
- b Under **Connect Workload Domains...**, click **Connect**.

The **Connect to Workload Domains** wizard opens to the **Modify Connection** page. This page lists all currently configured workload domains and enables you to connect vRealize Automation to each.

Important If you enable a connection between vRealize Automation and a workload domain, and then complete the **Connect to Workload Domains** wizard, you cannot disable the connection.

4 Select **Enable** for the desired workload domains.

5 If prompted, provide the Active Directory credentials used during the deployment of vRealize Automation and click **Next**. See [Deploy vRealize Automation in Cloud Foundation](#).

6 Review the connection and click **Finish**.

7 (Optional) Confirm the modified connection in vRealize Operations or vRealize Automation.

- a On the **vRealize Operations** or **vRealize Automation** page, click the product name link below the page title.

The vRealize Operations or vRealize Automation administrative opens to the Home page.

- b For vRealize Operations, navigate to **Administration > Solutions**.

The **Solutions** page displays the status of adapters for solutions connected to vRealize Operations. When successfully connected, the status indicates **Data Receiving**.

Note You may need to refresh the **Solutions** page to update the status.

Connect Workload Domains to vRealize Suite Products in Cloud Foundation

You can connect workload domains to vRealize Operations and vRealize Automation product deployments in Cloud Foundation using the SDDC Manager user interface.

Note This version of Cloud Foundation does not support connecting NSX-T workload domains to vRealize Automation through the SDDC Manager Dashboard. Use the vRealize Automation console instead.

Prerequisites

- Before you can connect the management domain or workload domains to vRealize Operations, it must be deployed. For more information, see [Deploy vRealize Operations in Cloud Foundation](#).
- Before you can connect workload domains to vRealize Automation, it must be deployed. For more information, see [Deploy vRealize Automation in Cloud Foundation](#).

Procedure

- 1 On the **SDDC Manager Dashboard**, navigate to **Inventory > Workload Domains**.

The **Workload Domains** page displays information for all workload domains.

- 2 Select the **Security** tab.

- 3 Click **Connect to vRealize Products**.

The **Connect to vRealize Products** wizard opens to the **Modify Connection** page. This page lists all currently configured workload domains and enables you to connect workload domains to either your vRealize Operations and vRealize Automation deployments.

- 4 Select **Enable** for the desired workload domain.

- 5 If prompted, provide the Active Directory credentials used during the deployment of vRealize Automation and click **Next**. See [Deploy vRealize Automation in Cloud Foundation](#).

- 6 Review the connection and click **Finish**.

Important If you enable a connection between vRealize Automation and a workload domain, and then complete the **Connect to Workload Domains** wizard, you cannot disable the connection.

Enable vRealize Log Insight in Cloud Foundation

You can connect vRealize Log Insight in Cloud Foundation to all workload domains in the SDDC Manager user interface.

Once enabled, you cannot disable the connection to vRealize Log Insight. All subsequently created workload domains will automatically connect and send logs to the vRealize Log Insight cluster.

Prerequisites

- Verify you have a valid license key for vRealize Log Insight, which is purchased separately from Cloud Foundation.

You can view your license in the vRealize Log Insight interface by navigating to **Management > License**.

- Verify that the vRealize Log Insight cluster is online and operational.

Procedure

- 1 On the **SDDC Manager Dashboard**, navigate to **Administration > vRealize Suite**.

The **vRealize Suite** navigation appears, listing the vRealize Suite products available for your Cloud Foundation system deployment.

- 2 Click **vRealize Log Insight**.

The **vRealize Log Insight** page displays.

The top portion of the page allows you to enable log collection for all workload domains. If not enabled, the **Enable** button is active.

The lower portion of the page displays the configuration details, including load balancer hostname, node size, and node count.

- 3 Click **Enable**.

After a moment, the page will update with a message indicating Connect Workload Domains to vRealize Log Insight in Progress. In **Tasks**, monitor the **Status** of the **Connect Workload Domains to vRealize Log Insight** action. Once Successful, vRealize Log Insight will collect logs from both the management workload domain and all additional workload domains.

Add a Node to a vRealize Operations Analytics Cluster

When your vRealize Operations analytics cluster needs more resources, you can add additional nodes.

You can use this procedure to expand the size of the vRealize Operations analytics cluster one node at a time. The original sizing configuration of your vRealize Operations deployment determines the total number of nodes that you can add. See the [vRealize Sizing Tool](#).

Prerequisites

- Forward and reverse DNS resolution is working for the FQDN of the new node.
- The IP address for the new node is in the same subnet as the existing vRealize Operations nodes.
- The datastore has enough disk space for the new node.

Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > vRealize Suite**.
- 2 Click **vRealize Operations**.
- 3 Click **Expand**.
- 4 Enter information for the new node and click **Next**.
 - Select **Data Node** for the component type.

- Enter the FQDN for the new node.
- By default, the new node uses the master node's password. If you want to use a different password for the new node, deselect the check box and enter a new password.

5 Click **Finish**.

Cloud Foundation validates the information and deploys the new node. If there are validation errors, fix the issues and click **Finish** to try again.

6 Monitor the deployment in the **Tasks** pane.

If the task fails, you can click **Restart Task** to try again, or click **Rollback** from **Administration > vRealize Suite > vRealize Operations** to return to a known good state. You can always restart a failed task. Depending on when a failure occurs, you may or may not be able to roll back the expansion. For rollback troubleshooting, see [KB 70991](#) and [KB 70992](#).

What to do next

If you replaced the vRealize Operations SSL certificate and configured the load balancer for SSL passthrough, see [Update Apache Configuration for New Nodes](#).

Update Apache Configuration for New Nodes

If you replaced the vRealize Operations SSL certificate and configured the load balancer for SSL passthrough, the you must configure Apache on each new node that you add to the vRealize Operations analytics cluster.

Procedure

- 1 Log into the new vRealize Operations Manager data node as root via SSH or Console.
- 2 Open `/usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf` in a text editor.
- 3 Find the `ServerName ${VCOPS_APACHE_SERVER_NAME}` line and insert a new line after it.
- 4 On the new line enter the following:

```
ServerAlias vrops-lb.vrack.vsphere.local vrops-master.vrack.vsphere.local
```

Replace `vrops-lb.vrack.vsphere.local` with the FQDN of vRealize Operations Manager load balancer and replace `vrops-master.vrack.vsphere.local` with the FQDN of the newly added vRealize Operations Manager data node.

- 5 Save and close the file.
- 6 Restart the apache2 service:

```
service apache2 restart
```

Downloading an Install Bundle

13

Cloud Foundation includes the following install bundles.

- VI workload domain install bundle is used to deploy later versions of the software components instead of the versions in your original Cloud Foundation installation. It includes software bits for vCenter Server and NSX for vSphere.
- vRealize Suite Lifecycle Manager install bundle is used to deploy vRealize Suite Lifecycle Manager.

This section describes the procedure for downloading install bundles from SDDC Manager. To download install bundles via a proxy server or in an offline mode, refer to the Download Bundles section in the *Cloud Foundation Upgrade Guide*.

Procedure

- 1 Log in to your My VMware Account.
 - a On the SDDC Manager Dashboard, click **Administration > Repository Settings**.
 - b Click **Authenticate**.
 - c Type your user name and password.
 - d Click **Authorize**.
- 2 On the SDDC Manager Dashboard, click **Repository > Bundles** in the left navigation pane. All available bundles are displayed. Install bundles display an Install Only Bundle label.
- 3 For the bundle you want to download, do one of the following:
 - Click **Download Now**.
The bundle download begins right away.
 - Click **Schedule Download**.
Select the date and time for the bundle download and click **Schedule**.
- 4 Navigate to **Repository > Download History > to see the downloaded bundles**.

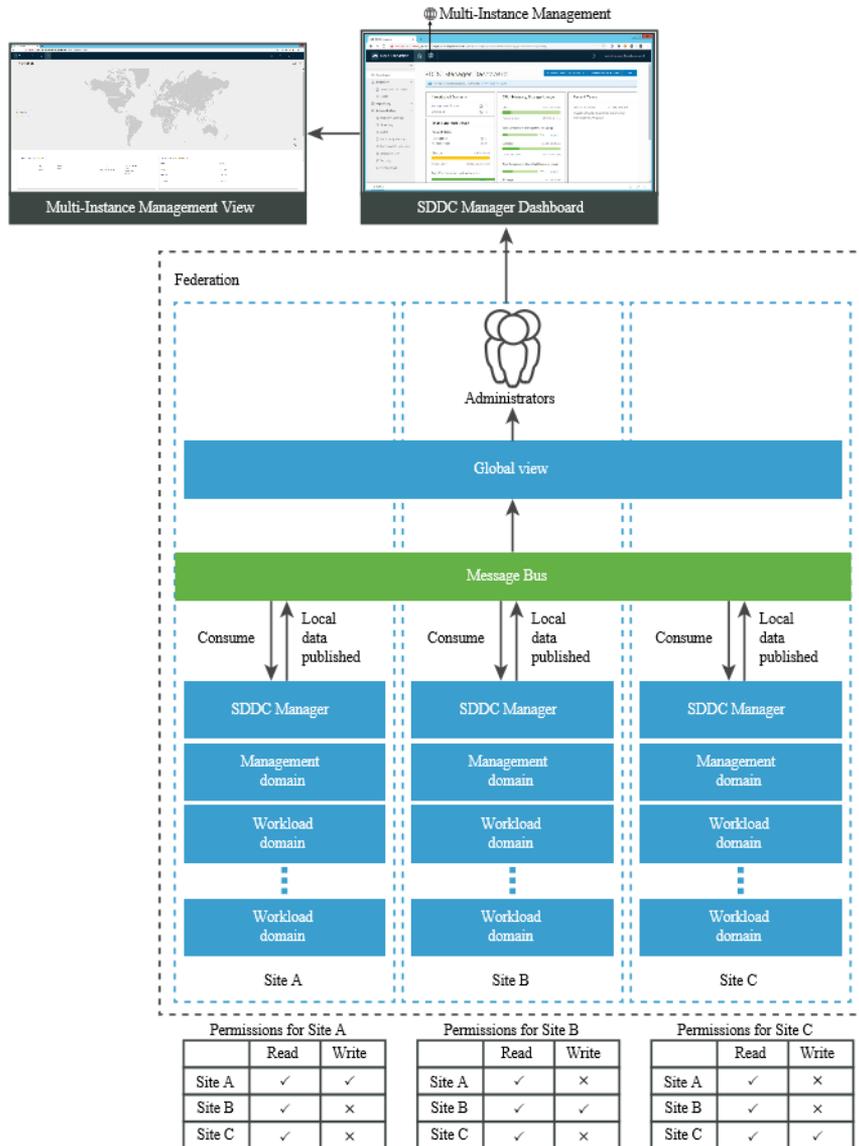
Multi-Instance Management

14

With the Multi-Instance Management feature, you can monitor and manage multiple Cloud Foundation instances from a single console.

Multiple Cloud Foundation instances can be managed together by grouping them into a federation, such that each member can view information about the entire federation and the individual instances within it. Federation members can view inventory across the Cloud Foundation instances in the federation as well as the available and used capacity (CPU, memory, and storage). This allows you to maintain control over the different sites and ensure that they are operating with the right degree of freedom and meeting compliance regulations for your industry. It also simplifies patch management by showing the number of patches available across sites in the global view.

Federation members communicate with each other via a message bus. Each participant publishes their local data to the message bus and the remaining participants can read this data for global visibility across the federation.



An instance can see details about the federation only if it is a member of the federation, and can belong only to a single federation at a time. It is possible to create multiple federations within an organization; however, there is no global visibility between federations. For example, it might be desirable to have a dev-test federation and a production federation. In such an example, members of dev-test can see other dev-test members but they are not able to see production members.

Federation members can either be controllers or regular members. A controller member has capabilities of a regular member and runs some additional message bus components to allow multi-instance management to work.

A controller member can invite other instances to become members as controller or regular members. The controller role can be granted to a maximum of three instances within a federation. High Availability of multi-instance management functionality is only possible when there are exactly three controllers in the federation. Though the controller members can be at

any three sites in the federation, it is recommended that each controller is in a different availability zone. The instance who created the federation is automatically granted the controller role. If you only have two instances in the federation, there is no need to create both as controllers. Multi-Instance management works with two Cloud Foundation sites; however, if one fails then the multi-instance capability is not available on the other site.

This chapter includes the following topics:

- [About the Multi-Instance Management Dashboard](#)
- [Create a Federation](#)
- [Invite a Cloud Foundation Instance to Join a Federation](#)
- [Join a Federation](#)
- [Leave a Federation](#)
- [Dismantle a Federation](#)

About the Multi-Instance Management Dashboard

The Multi-Instance Management Dashboard displays the inventory and capacity across the federation.

You access the Multi-Instance Management Dashboard by clicking the Multi-Instance View icon in the top left corner of the SDDC Manager Dashboard.

The screenshot shows the VMware Cloud Foundation SDDC Manager Dashboard. A red arrow points to the 'Multi-Instance Management icon' in the top navigation bar. The dashboard displays the following information:

- 4 Workload Domains:** Management Domain (1 icon), VI Domain (3 icons).
- Host Type and Usage:**
 - Host Types: Hybrid Host (7 icons), All Flash Host (2 icons).
 - Usage: 9 Hosts Total (9 Hosts Used, 0 Hosts Unallocated).
- CPU, Memory, Storage Usage:**
 - CPU: 143.64 GHz Total (14.55 GHz Used, 129.09 GHz Free).
 - Memory: 531.25 GB Total (254.6 GB Used, 266.65 GB Free).
 - Top Domains in allocated CPU Usage: Vi-NEXT (17%), MGMT (11%), CNF, Delph V1.1 (6%).
 - Top Domains in allocated Memory Usage: (Data partially obscured).
- Recent tasks:**
 - 11/13/17, 6:11 AM: Succeeded (Page 1 Create new Virtual Infrastructure domain).
 - 10/30/17, 3:06 PM: Running (CMS).
 - 10/30/17, 9:27 AM: Succeeded (Bringup).
 - 8/1/17, 4:28 PM: Pending (CMS).
 - 8/1/17, 4:42 PM: MANAGEMENT.

Before a federation is created, the dashboard displays a create and join option.

Welcome to Multi-Instance Management

Please select one of the following based on your role.

Create a Federation



Only 1 user should create a federation for a given organization. This instance will become the first Controller. ⓘ

CREATE

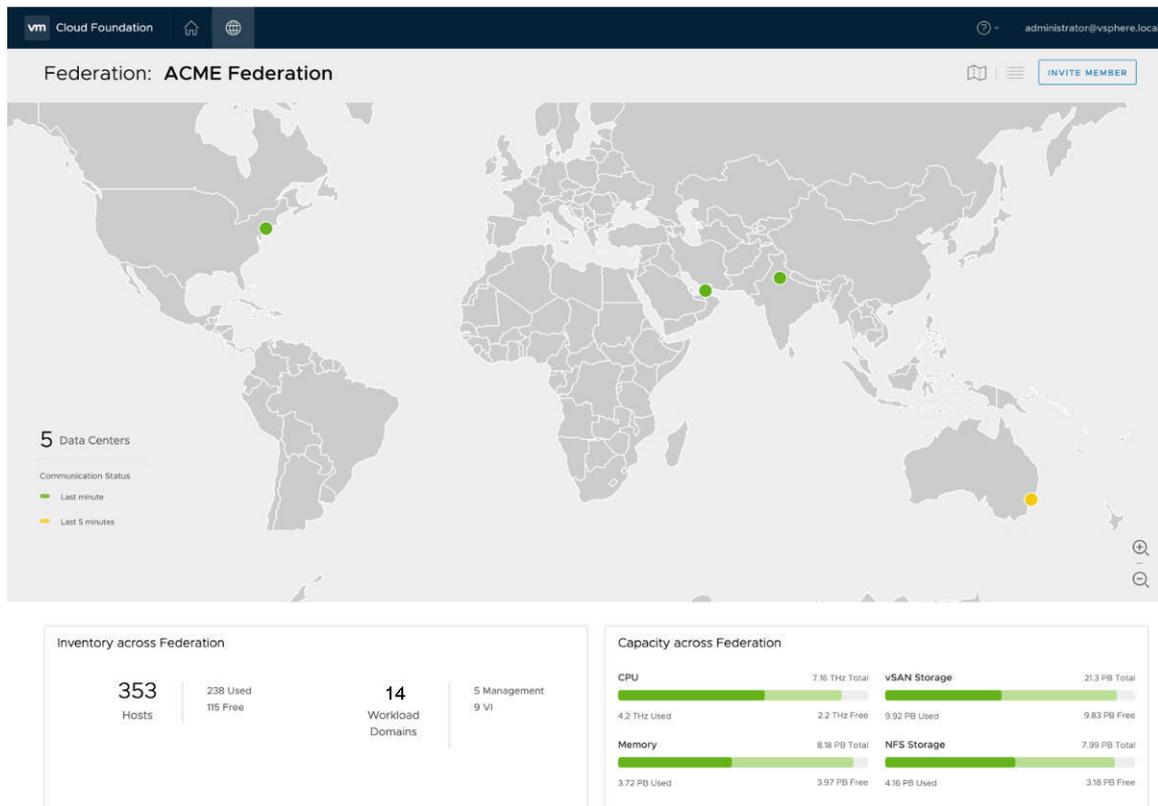
Join a Federation



Most users will be joining an established federation by invitation. Please refer to the instructions you received in order to join.

JOIN

After a federation is created, the Multi-Instance Management Dashboard displays a world map showing the federation members as dots on the map.



The dot color depends on the communication status between the federation members - green if they communicated within the last two minutes, yellow if they communicated within the previous five minutes, and red if they have not communicated for more than ten minutes. You can see the following information here:

- Hover over a dot to see the member name and location.



- Click the member location dot to open a panel on the right side with detailed information about the Cloud Foundation instance. The panel also displays available software updates.

Palo Alto ● Controller ⓘ

Data Center

Location: Palo Alto, United States

Hosts: 7
4 Used | 3 Free

Workload Domains: 1
1 Management | 0 VI

Capacity:

CPU 76.61 GHz Total

11.2 GHz Used 65.41 GHz Free

Memory 468.75 GB Total

132.49 GB Used 336.26 GB Free

vSAN Storage 4.47 TB Total

0.63 TB Used 3.84 TB Free

Software Updates: ⓘ

The Inventory section in the bottom half of the dashboard displays the number of hosts and workload domains along with a breakdown of the workload domain type. The capacity section displays the used and available CPU, memory, and storage across the federation.

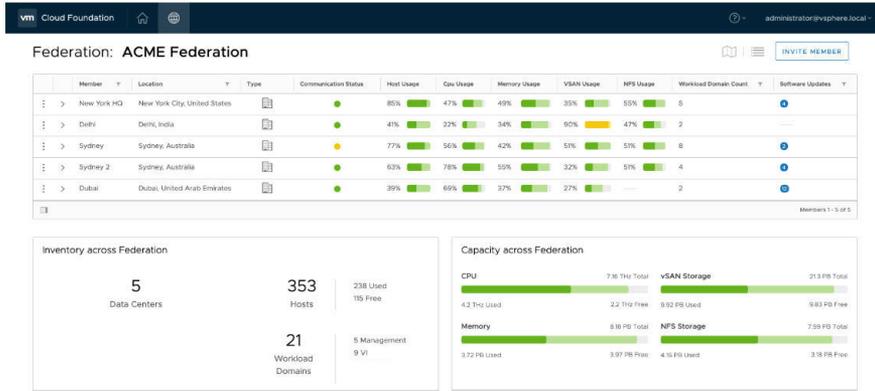
Inventory across Federation

353 Hosts	238 Used 115 Free	21 Workload Domains	5 Management 9 VI
---------------------	----------------------	----------------------------------	----------------------

Capacity across Federation

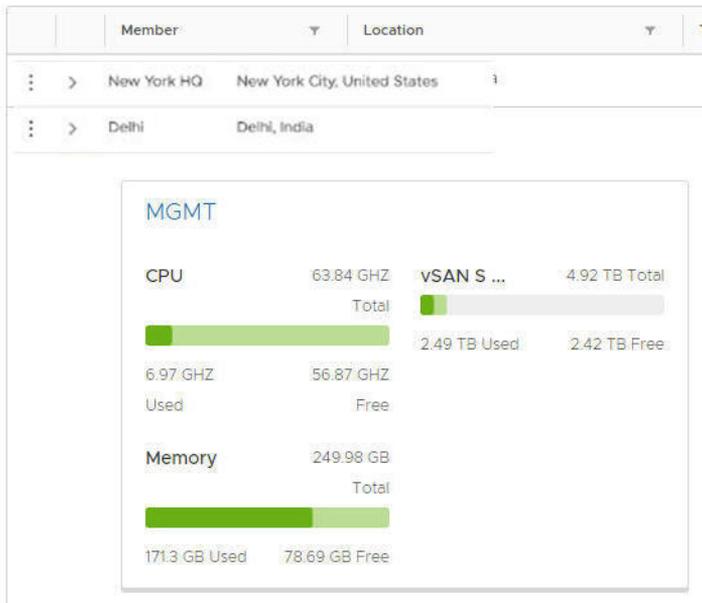
CPU 7.16 THz Total	vSAN Storage 21.3 PB Total
4.2 THz Used 2.2 THz Free	9.92 PB Used 9.83 PB Free
Memory 8.18 PB Total	NFS Storage 7.99 PB Total
3.72 PB Used 3.97 PB Free	4.16 PB Used 3.18 PB Free

Click the Table icon at the top right of the dashboard to display member information in a grid format. Information for all federation members is displayed in a tabular format.



Clicking the arrow (➤) next to the member name displays the CPU, memory, and storage usage for that member.

Federation: ACME Federation



Clicking MGMT takes you to the management domain of that member.

Create a Federation

A federation is a group of Cloud Foundation instances, such that each member can view information about the other Cloud Foundation instances in the group. The federation creator is granted the controller role by default.

You can create multiple federations within your organization, but global visibility is available only within a federation. Members can belong to only a single federation at a time.

See [VMware Configuration Maximums](#) for information about the maximum number of SDDC Manager instances that can be managed using Multi-Instance Management

Prerequisites

- The SDDC Manager certificate requires TLS Web server authentication and TLS Web client authentication extended key usage (EKU). If you are using Microsoft CA or a third-party CA, you must ensure that the Certificate Authority template is configured with both these EKUs.
- Retrieve the FQDN of your SDDC Manager.

Procedure

- 1 On the SDDC Manager dashboard, click the Multi-Instance View () icon at the top of the window.
- 2 Click **Create**.

Create Federation

Establishing a federation requires two steps:

1. Register the first member (who is designated as a Controller).
2. Invite additional members to the federation.

Responsibilities of the Controller:

- Only Controllers may invite new members to a federation.
- Controllers can only dismantle a federation once all its members have left the federation.
- Providing high availability.

Federation Name	<input type="text" value="ACME Federation"/>
Member Name 	<input type="text" value="Palo Alto"/>
FQDN	<input type="text" value="sddc-manager.vrack.vsphere.local"/>
Country	<input type="text" value="United States"/> 
City	<input type="text" value="Palo Alto"/> 

- 3 Enter a name for the federation.
- 4 Enter a display name for the member. You may want to base this on the location of this Cloud Foundation instance.
- 5 Type the FQDN of the SDDC Manager.
- 6 Select the city and country of this Cloud Foundation instance and click **Create**.

Results

It can take a few minutes for the federation to be created. After the federation is created, the Multi Instance Management Dashboard is displayed. The federation location is marked with a green dot on the world map. You can zoom in or out of the map.

The dashboard also displays the inventory (hosts and workload domains) and capacity (CPU, memory, and storage) across the federation. These details are updated when additional members join the federation.

What to do next

Invite a Cloud Foundation instance to join the federation.

Invite a Cloud Foundation Instance to Join a Federation

You can invite Cloud Foundation instances to join a federation. They can be invited as a controller or a regular member. High Availability of multi-instance management functionality is only possible when there are exactly three controllers in the federation. Though the controller members can be at any three sites in the federation, it is recommended that each controller is in a different availability zone.

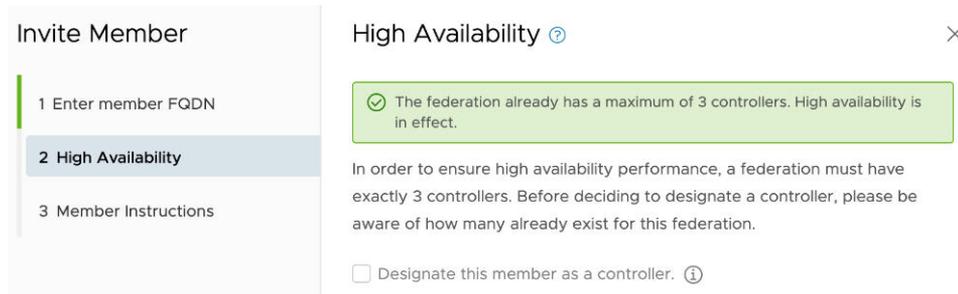
Prerequisites

You must be a controller in the federation and have the FQDN of the member you are inviting.

Procedure

- 1 On the top right corner of the Multi-Instance Management dashboard, click **Invite Member**.
- 2 Enter the SDDC Manager FQDN of the member you are inviting and click **Check Certificate**.
The invited member's certificate thumbprint is displayed.
- 3 Validate the thumbprint and click **Confirm fingerprint**.
- 4 Click **Next**.
- 5 Select the check box on the High Availability page if you want to designate the controller role to the member.
- 6 Click **Next**.

The Instructions page displays the URL that the invited member needs to access.



- 7 Click **Copy Info** to copy the information displayed on this page or copy the URL manually and send it to the member through an offline method.

What to do next

The invitation and joining process is a coordinated effort between the invitee controller and joining member. An additional dot on your Multi-Instance Management Dashboard indicates that the member you invited is joining the federation. When a controller joins a federation, it can take a few minutes for the federation to stabilize.

Join a Federation

You can join a federation as a controller or member depending on the assigned role in the invitation. An invitation is valid for ten days. You must request a new invitation after this period. If a new invitation is generated for the same site, only the latest invite is valid.

Prerequisites

- Verify that the federation is healthy - there should be no red dots on the world map of the Multi-Instance Management Dashboard. A red dot indicates that SDDC Manager is unable to communicate with that member. A controller must remove the member using the leave API. See the *VMware Cloud Foundation 4.0 API Reference Guide*.
- The SDDC Manager certificate requires TLS Web server authentication and TLS Web client authentication extended key usage (EKU). If you are using Microsoft CA or a third-party CA, you must ensure that the Certificate Authority template is configured with both these EKUs.
- Retrieve the FQDN of your SDDC Manager.

Join a Federation by Clicking an Invitation

If you join a federation by clicking the invitation you received, federation details are pre-populated in the UI.

Prerequisites

Retrieve the invitation you received.

Procedure

- 1 Click the URL in the invitation you received.

The Join Federation window displays the role assigned in the invitation, the FQDN of the invited member, token, and the FQDN of the controller member who invited you to join the federation.

Join Federation

✔ Certificate is validated successfully.
✕

Member Name ⓘ	<input type="text" value="Delhi"/>
Member Role	<input type="text" value="Member"/>
FQDN ⓘ	<input type="text" value="delhi.mydomain.local"/>
Country	<input type="text" value="India"/>
City	<input type="text" value="Delhi"/>
Token ⓘ	<input type="text" value="jhdjfJHJGDJKKDF57642j4JHBBdkjndnk"/>
FQDN of Controller ⓘ	<input type="text" value="newyork.mydomain.local"/>

2 Click **Join**.

Results

The join process is initiated. If you see an error, resolve the issue and then request a new invitation and follow the steps described above. After the join process is successful, your Cloud Foundation instance becomes a member of the federation.

After a controller joins or leaves a federation, Kafka is restarted on all controllers in the federation. It can take several minutes for the federation to stabilize even after the dashboard is refreshed. If an operation performed on the dashboard during this time fails, re-try the operation.

Join a Federation through the Multi-Instance Management Dashboard

You can join a federation through the Multi-Instance Management Dashboard.

Procedure

- 1 Click **Join Federation** on the Multi-Instance Management Dashboard.
- 2 Type a display name for the site to be added.
- 3 Select the member role as indicated in the invitation you received.
- 4 Type the FQDN of your SDDC Manager.
- 5 Select the country and city for your site.
- 6 Type the token as indicated in the invitation you received.
- 7 Type the FQD of the controller who invited you.
- 8 Click **Join**.

Results

The join process is initiated. If you see an error, resolve the issue and then request a new invitation and follow the steps described above. After the join process is successful, your Cloud Foundation instance becomes a member of the federation.

After a controller joins or leaves a federation, Kafka is restarted on all controllers in the federation. It can take several minutes for the federation to stabilize even after the dashboard is refreshed. If an operation performed on the dashboard during this time fails, re-try the operation.

Leave a Federation

Leaving a federation removes the Multi-Instance Management view from your SDDC Manager Dashboard.

If you are a controller, you can leave a federation only if there is at least one more controller in the federation. If you are the only controller member in a federation, you must dismantle a federation instead of leaving it.

Prerequisites

Verify that the federation is healthy - there should be no red dots on the world map of the Multi-Instance Management Dashboard. A red dot indicates that SDDC Manager is unable to communicate with that member. A controller must remove the member using the leave API. See the *VMware Cloud Foundation 4.0 API Reference Guide*.

Procedure

- 1 Click the Grid icon at the top right of the Multi-Instance Management Dashboard.
- 2 In the member table, click the dot icon next to your member name and click **Leave Federation**.
- 3 Type the federation name and click **Leave**.

What to do next

Do not perform any operation for a few minutes after leaving a federation.

Dismantle a Federation

You can dismantle a federation if you are the last controller member in the federation. Only members with the controller role can dismantle a federation.

Procedure

- 1 Click the Grid icon at the top right of the Multi-Instance Management Dashboard.
- 2 In the member table, click the dot icon next to your member name and click **Dismantle Federation**.
- 3 Type the federation name and click **Dismantle**.

What to do next

After the federation is dismantled, the Create Federation screen is displayed instead of the Multi-Instance Management Dashboard.

Stretching Clusters

15

You can stretch a cluster in the management domain or in a VI workload domain across two availability zones.

This section describes how to stretch an NSX for vSphere cluster. To stretch an NSX-T cluster, see [Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation](#).

You may want to stretch a cluster for the following reasons.

- **Planned maintenance**

You can perform a planned maintenance on an availability zone without any downtime and then migrate the applications after the maintenance is completed.

- **Automated recovery**

Stretching a cluster automatically initiates VM restart and recovery, and has a low recovery time objective for the majority of unplanned failures.

- **Disaster avoidance**

With a stretched cluster, you can prevent service outages before an impending disaster such as a hurricane or rising flood levels.

This chapter includes the following topics:

- [About Availability Zones and Regions](#)
- [Prerequisites for Stretching a Cluster](#)
- [Stretch a Cluster](#)
- [Unstretch a Cluster](#)
- [Expand a Stretched Cluster](#)
- [Replace a Failed Host in a Stretched Cluster](#)

About Availability Zones and Regions

An availability zone is a collection of infrastructure components. Each availability zone is isolated from other availability zones to prevent the propagation of failure or outage across the data center. An availability zone runs on its own physically distinct, independent infrastructure, and is

engineered to be highly reliable. Each zone should have independent power, cooling, network, and security. Additionally, these zones should be physically separate so that even uncommon disasters affect only one zone.

The physical distance between availability zones is short enough to offer low, single-digit latency (less than 5 ms) and large bandwidth (10 Gbps) between the zones. Hence, availability zones can either be two distinct data centers in a metro distance, or two safety or fire sectors (data halls) in the same large-scale data center.

The recommended minimum number of hosts in each availability zone is 4 hosts and the maximum is 15 hosts. If you are expanding a cluster, you must add hosts in pairs. Each host in the pair must have the same CPU, memory, and storage.

A region is a Cloud Foundation instance.

Note Cloud Foundation supports stretching a cluster across two availability zone within a region.

Prerequisites for Stretching a Cluster

Before you can stretch a cluster, you must meet the prerequisites.

- Open the [Deployment of Multiple Availability Zones](#) document and read it to understand the requirements.
- Ensure a vSAN Enterprise license has been applied to the cluster to be stretched. A vSAN Enterprise license is required for stretching a cluster. See [KB 70328](#) for information about a known licensing issue.
- The management VLAN between the two availability zones must be stretched.
- All VMs on an external network must be on a virtual wire. If they are on a VLAN, that VLAN must be stretched as well.
- Each availability zone must have its own vMotion, vSAN, and VXLAN networks.
- The vMotion, vSAN, and VXLAN networks require L3 routing between the availability zones. vSAN networks must also have L3 routing to the vSAN network of the witness host.
- The VLAN IDs must be identical on both availability zones.
- Each stretched cluster requires a vSAN witness appliance in a third party location. The witness appliance should be running the same version of ESXi as the ESXi hosts in the stretched cluster. The maximum RTT on the witness is 200ms. Follow the steps described in [Deploy the vSAN Witness Host in Region B](#) to add and configure a vSAN witness.
- If you are stretching a cluster in a VI workload domain, you must stretch the management domain cluster first. vCenter Servers for all workload domains are in the management domain. Hence, you must protect the management domain to ensure that you can access and manage the workload domains.

- Ensure that you have enough hosts such that there is an equal number of hosts on each availability zone. This is to ensure that there are sufficient resources in case an availability zone goes down completely.
- TCP port and UDP Ports needs to be open for witness traffic between the witness host and the vSAN cluster data nodes. See KB article [52959](#).

Cloud Foundation Networks

Network Name	Connectivity to AZ2	Minimum MTU	Maximum MTU
vSAN	L3	1500	9000
vMotion	L3	1500	9000
VXLAN (VTEP)	L3	1600	9000
Management	L2	1500	9000
Witness Management	L3	1500	9000
Witness vSAN	L3	1500	9000

Stretch a Cluster

This procedure describes how to stretch a cluster across two availability zones.

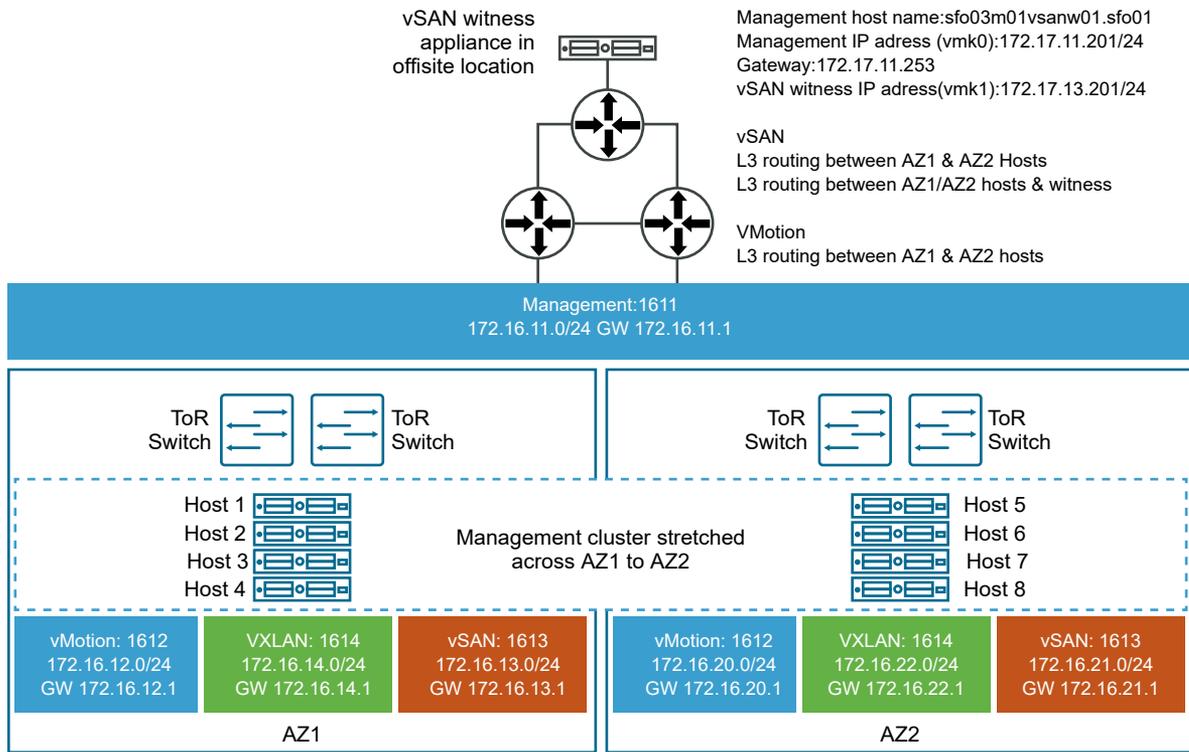
As an example, we will follow a use case with two availability zones in two buildings in an office campus - AZ1 and AZ2. Each availability zone has its own power supply and network. The management domain is on AZ1 and contains the default cluster, SDDC-Cluster1. This cluster contains four ESXi hosts. AZ1 also contains the default bring-up pool, bringup-networkpool.

vSAN network	VLANID=1613
	MTU=9000
	Network=172.16.13.0
	netmask 255.255.255.0
	gateway 172.16.13.1
	IP range=172.16.13.11 - 172.16.13.59
vMotion network	VLANID=1612
	MTU=9000
	Network=172.16.12.0
	netmask 255.255.255.0
	gateway 172.16.12.1
	IP range=172.16.12.11 - 172.16.12.59

There are four ESXi hosts in AZ2 that are not in the Cloud Foundation inventory yet.

We will stretch the default cluster SDDC-Cluster1 in the management domain from AZ1 to AZ2.

Figure 15-1. Stretch Cluster Example



Prerequisites

Ensure you have completed the steps listed in [Prerequisites for Stretching a Cluster](#).

Procedure

- 1 Create a network pool, AZ2-networkpool, on AZ2. See [Create a Network Pool](#).

Based on our example, here are the network details for the network pool.

vSAN network	VLANID=1613
	MTU=9000
	Network=172.16.21.0
	Netmask=255.255.255.0
	Gateway=172.16.21.1
	IP range= 172.16.21.11 - 172.16.21.59
vMotion network	VLANID=1612
	MTU=9000
	Network=172.16.20.0
	netmask 255.255.255.0
	gateway 172.16.20.1
	IP range= 172.16.20.11 - 172.16.20.59

- 2 Commission the four hosts in AZ2 and associate them with AZ2-networkpool. In our example, these are 172.16.11.105, 172.16.11.106, 172.16.11.107, 172.16.11.108.

See [Commission Hosts](#).

- 3 Retrieve the FQDNs of the hosts in AZ2.
 - a On the SDDC Manager Dashboard, click **Hosts**.
 - b Note down the FQDNs for the hosts in AZ2.
- 4 SSH in to the SDDC Manager VM using the **vcf** administrative user account.
- 5 Enter **su** to switch to the root user and navigate to the `/opt/vmware/sddc-support` directory.
- 6 Enter the following command:

```
./sos --stretch-vsan --sc-domain <DOMAIN NAME> --sc-cluster <CLUSTER NAME> --sc-hosts
<HOSTFQDN,HOSTDQND2,...> --witness-host-fqdn <WITNESS HOST FQDN> --witness-vsan-ip <WITNESS VSAN
IP> --witness-vsan-cidr <WITNESS VSAN CIDR> --esxi-license-key <LICENSE KEY>
```

Example input and response:

```
root@sddc-manager [ /opt/vmware/sddc-support ]# ./sos --stretch-vsan --sc-domain MGMT --sc-
cluster SDDC-Cluster1 --sc-hosts esxi-5.vrack.vsphere.local,esxi-6.vrack.vsphere.local --witness-
host-fqdn esxi-11.vrack.vsphere.local --witness-vsan-ip 10.0.12.96 --witness-vsan-cidr
10.0.12.0/24 --esxi-license-key AAAAA-BBBBB-CCCCC-DDDDD-EEEE
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-02-13-08-51-34-12479
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-02-13-08-51-34-12479/sos.log
Starting vSAN stretched cluster operations..
Api Response:{
  "taskId": "d670ff00-24a3-4739-b5ff-b5317d709f36",
  "resourceId": "0f8d112d-aa3f-4ca8-a8bd-b95e0e1ea2f5",
  "resourceType": "ESXI",
  "state": "IN_PROGRESS",
  "description": "Stretch vSAN Cluster - SDDC-Cluster1 in VMware Cloud Foundation",
  "timestamp": 1550047894872,
  "id": "d670ff00-24a3-4739-b5ff-b5317d709f36"
}
```

- 7 Monitor the state of the task in the SDDC Manager Dashboard.

- 8 When the task completes successfully, validate that stretched cluster operations are working correctly by logging in to the vSphere Web Client.

The time it takes to complete a policy compliance check depends on the number of VMs in the cluster.

- a Check the vSAN Health page.
 - 1 On the home page, click **Host and Clusters** and then select the stretched cluster (SDDC-Cluster1 in our example).
 - 2 Click **Monitor > vSAN > Health**.
 - 3 Click Retest.
 - 4 Fix errors, if any.
- b Check the vSAN Storage Policy page.
 - 1 On the home page, click **Policies and Profiles > VM Storage Policies > vSAN Default Storage Policies**.
 - 2 Select the policy associated with the vCenter Server for the stretched cluster.
 - 3 Click **Monitor > VMs and Virtual Disks**.
 - 4 Click **Refresh**.
 - 5 Click **Trigger VM storage policy compliance check**.
 - 6 Check the **Compliance Status** column for each VM component.
 - 7 Fix errors, if any.

What to do next

It is recommended that you add new VMs to the primary AZ and associate the host rule for the new VM with the primary AZ rule.

Unstretch a Cluster

This procedure describes how to unstretch a vSAN cluster which is stretched across two availability zones and convert it to a standard vSAN cluster.

As an example, we will consider a use case in which there two availability zones, AZ1 and AZ2, in two buildings in an office campus. Each availability zone has its own power supply and network. The management domain is on AZ1 and contains the default cluster, SDDC-Cluster1, which is stretched between AZ1 and AZ2. There are four ESXi hosts each in AZ1 and AZ2, which are categorized into the primary fault domain and secondary fault domain. This example unstretches the default cluster in the management domain and converts the stretched cluster to standard vSAN cluster.

Prerequisites

You must have a stretched cluster.

Procedure

- 1 SSH in to the SDDC Manager VM using the **vcf** administrative user account.
- 2 Enter **su** to switch to the root user and navigate to the `/opt/vmware/sddc-support` directory.
- 3 Enter the following command:

```
./sos --unstretch-vsan --sc-cluster <CLUSTER NAME> --sc-domain <DOMAIN NAME>
```

Example input and response:

```
root@sddc-manager [ /opt/vmware/sddc-support ]# ./sos --unstretch-vsan --sc-cluster SDDC-Cluster1
--sc-domain MGMT
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-02-21-07-36-18-66388
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-02-21-07-36-18-66388/sos.log
Starting vSAN stretched cluster operations..
vSAN un-stretch operation started..
Api Response:{
  "taskId": "9c3b0975-be3c-42f2-8d1a-3d708c2c3263",
  "resourceId": "5f6cac74-6fbb-4570-b240-1a0ed5a54118",
  "resourceType": "ESXI",
  "state": "IN_PROGRESS",
  "description": "Un-Stretch vSAN Stretched Cluster - SDDC-Cluster1 in VMware Cloud Foundation",
  "timestamp": 1550734579412,
  "id": "9c3b0975-be3c-42f2-8d1a-3d708c2c3263"
}
```

- 4 Monitor the state of the task in the SDDC Manager Dashboard.

When the task completes, all the hosts from the AZ2 are removed from the cluster and cluster is converted to standard vSAN cluster.

- 5 Validate that vSAN cluster operations are working correctly by logging in to the vSphere Web Client.
 - a Check the vSAN Health page.
 - 1 On the home page, click **Host and Clusters** and then select the vSAN cluster (SDDC-Cluster1 in our example).
 - 2 Click **Monitor > vSAN > Health**.
 - 3 Click Retest.
 - 4 Fix errors, if any.
 - b Check the vSAN Storage Policy page.
 - 1 On the home page, click **Policies and Profiles > VM Storage Policies > vSAN Default Storage Policies**.
 - 2 Select the policy associated with the vCenter Server for the vSAN cluster.
 - 3 Click **Monitor > VMs and Virtual Disks**.
 - 4 Click **Refresh**.
 - 5 Click **Trigger VM storage policy compliance check**.
 - 6 Check the **Compliance Status** column for each VM component.
 - 7 Fix errors, if any.

Expand a Stretched Cluster

You can expand a stretched cluster by adding hosts. The hosts need to be added in pairs (such as two, four, six, or eight). Half of the hosts you add to the cluster will be added to the first availability zone and the other half will be added to the second availability zone.

Prerequisites

You must have a stretched cluster and available hosts.

Procedure

- 1 Commission the additional hosts to Cloud Foundation. For each pair of hosts, associate one with the network pool in AZ1 and the other with the network pool in AZ2.
See [Commission Hosts](#).
- 2 SSH in to the SDDC Manager VM using the **vcf** administrative user account.
- 3 Enter **su** to switch to the root user and navigate to the `/opt/vmware/sddc-support` directory.
- 4 Enter the following command:

```
./sos --expand-stretch-cluster --sc-domain <DOMAIN NAME> --sc-cluster <CLUSTER NAME> --sc-hosts <HOSTFQDN,HOSTDQND2,...> --esxi-license-key <LICENSE KEY>
```

Example input and response:

```

root@sddc-manager [ /opt/vmware/sddc-support ]# ./sos --expand-stretch-cluster --sc-domain MGMT --
sc-cluster SDDC-Cluster1 --sc-hosts esxi-9.vrack.vsphere.local, esxi-10.vrack.vsphere.local --
esxi-license-key AAAAA-BBBBB-CCCC-DDDD-EEEE
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-02-20-10-04-32-123007
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-02-20-10-04-32-123007/sos.log
Starting vSAN stretched cluster operations..
expand vSAN stretched cluster operation started
Api Response:{
  "taskId": "6e4b13d9-ead-408b-a595-4e89ef885a3e",
  "resourceId": "0c518498-b302-40ae-abc4-10addead7bc2",
  "resourceType": "ESXI",
  "state": "IN_PROGRESS",
  "description": "Prepare vSAN Cluster - SDDC-Cluster1 for Stretch in VMware Cloud Foundation",
  "timestamp": 1550657073615,
  "id": "6e4b13d9-ead-408b-a595-4e89ef885a3e"
}

```

- 5 Monitor the state of the task in the SDDC Manager Dashboard.
- 6 If required, SSH in to each newly added hosts and add a static route to the vSAN network of the witness host. Add static routes in the witness if it could not reach the vSAN network of the newly added hosts.
- 7 Move the added host to the appropriate availability zone so that the cluster is back to containing two fault domains again.
 - a On the SDDC Manager Dashboard, click **Inventory > Workload Domains** and then click **View Details**.
 - b Click the name of the domain containing the stretched cluster, for example, MGMT.
 - c Click the **Services** tab and click the vCenter Server launch icon and log in to the vSphere Web Client.
 - d In the vSphere Web Client, select the stretched cluster. Then select **Configure > vSAN > Fault Domains & Stretched Cluster**.
 - e Select the first newly added host associated with the network pool on AZ1 and drag it to AZ1.
 - f Select the second newly added host associated with the network pool on AZ2 and drag it to AZ2.
- 8 Add these hosts to the VMHost rule so that you can deploy VMs on all hosts.
 - a In the vSphere Web Client, select **Hosts and Clusters** and then select the stretched cluster.
 - b Select **Configure > VM/Host Rules**.

- c Select the `<cluster_name>` rule and click **Add**.
 - d Select the ESXi hosts newly added to availability zone 1 and click **OK**.
- 9** Update the value for **Host failure cluster tolerates** to the number of hosts in AZ1 after the expansion.
- a Log in to the management vCenter Server in the vSphere Web Client.
 - b From the **Home** menu, select **Hosts and Clusters** and expand the stretched cluster.
 - c Select the stretched cluster and click the **Configure** tab.
 - d Under **Services**, click **vSphere Availability**, and click **Edit**.
 - e On the Admission Control page of the Edit Cluster Settings dialog box, set **Host failures cluster tolerates** to the number of hosts in AZ1 and click **OK**.

Replace a Failed Host in a Stretched Cluster

If a host or host component in a stretched cluster fails, it is recommended that you replace the host with a new host.

Prerequisites

- Image the replacement host with the same ESXi version as the other hosts in the cluster.
- Check the health of the cluster.

See "Check vSAN Health" in *Administering VMware vSAN*.

Procedure

- 1** Remove the failed host from the cluster.
See [Remove a Host from a Cluster in a Workload Domain](#)
- 2** Decommission the host.
See [Decommission Hosts](#).
- 3** Commission the replacement host to the same network pool as the removed host.
See [Commission Hosts](#) .
- 4** SSH in to the SDDC Manager VM using the **vcf** administrative user account.
- 5** Enter **su** to switch to the root user and navigate to the `/opt/vmware/sddc-support` directory.
- 6** Enter the following command:

```
./sos --expand-stretch-cluster --sc-domain <DOMAIN NAME> --sc-cluster <CLUSTER NAME> --sc-hosts <REPLACEMENT HOSTFQDN> --esxi-license-key <LICENSE KEY>
```

Example input and response:

```
root@sddc-manager [ /opt/vmware/sddc-support ]# ./sos --expand-stretch-cluster --sc-domain MGMT --
sc-cluster SDDC-Cluster1 --sc-hosts esxi-11.vrack.vsphere.local --esxi-license-key AAAAA-BBBBB-
CCCC-DDDDD-EEEE
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-02-20-10-04-32-123007
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-02-20-10-04-32-123007/sos.log
Starting vSAN stretched cluster operations..
expand vSAN stretched cluster operation started
Api Response:{
  "taskId": "6e4b13d9-eead-408b-a595-4e89ef885a3e",
  "resourceId": "0c518498-b302-40ae-abc4-10addead7bc2",
  "resourceType": "ESXI",
  "state": "IN_PROGRESS",
  "description": "Prepare vSAN Cluster - SDDC-Cluster1 for Stretch in VMware Cloud Foundation",
  "timestamp": 1550657073615,
  "id": "6e4b13d9-eead-408b-a595-4e89ef885a3e"
}
```

- 7 Monitor the state of the task in the SDDC Manager Dashboard.
Wait until the task completes successfully.
- 8 If required, SSH in to the newly added host and add a static route to the vSAN network of the witness host. Add static routes in the witness if it could not reach the vSAN network of the newly added host.
- 9 In the vSphere Web Client, move the host to the appropriate availability zone.
 - a On the SDDC Manager Dashboard, click **Inventory > Workload Domains** and then click **View Details**.
 - b Click the name of the domain containing the stretched cluster, for example, MGMT.
 - c Click the **Services** tab and click the vCenter Server launch icon and log in to the vSphere Web Client.
 - d In the vSphere Web Client, select the stretched cluster. Then select **Configure > vSAN > Fault Domains & Stretched Cluster**.
 - e Select the newly added host and drag it to the appropriate availability zone.
- 10 If the host belongs to AZ1, add the host to the AZ1 VMHost rule. If the host belongs to AZ2, no operation is required.
 - a In the vSphere Web Client, select **Hosts and Clusters** and then select the stretched cluster.
 - b Select **Configure > VM/Host Rules**.
 - c Select the appropriate rule and click **Add**.
 - d Select the newly added host and click **OK**.

Results

vSAN automatically rebuilds the stretch cluster.

Composability Management

16

With composability, you can dynamically configure servers to meet the needs of your workloads without physically moving any hardware components. You bind disaggregated hardware components (compute, network, storage, and offload components) together to create a logical system based on the needs of your applications. These logical systems function like traditional rack mount systems.

The Cloud Foundation composability feature is available for HPE Synergy and Dell MX servers and uses the Redfish translation layer to connect to the composable hardware infrastructure. Redfish Translation Layer supports data models used to get composable resources and zones restrictions from the hardware infrastructure. It is designed to be extensible and vendor agnostic. You must obtain and install the Redfish appliance from the composable hardware vendor.

It is recommended that you compose and decompose servers only through Cloud Foundation, and not through the vendor software.

Note Ensure that you follow the restriction on storage sled placement on Dell MX servers. Refer to vendor documentation for more information.

This chapter includes the following topics:

- [Configure Translation Layer](#)
- [Compose a Server](#)
- [View Composability Information](#)
- [Add Storage](#)
- [Remove Storage](#)
- [Decompose a Server](#)

Configure Translation Layer

Redfish translation layer is the interface between SDDC Manager and hardware vendor. You must configure this translation layer by providing the Redfish translation layer URL and credentials.

Procedure

- 1 As a best practice, increase the queue capacity for the thread pool.
 - a Open the `application-prod.properties` file:


```
vi /opt/vmware/vcf/operationsmanager/config/application-prod.properties
```
 - b Update the queue capacity line as follows:


```
om.executor.queuecapacity=300
```
 - c Save and close the file.

- 2 If you are using a self-signed certificate, import the Redfish certificate from the Redfish VM to SDDC Manager VM by following the steps below. If you are using a CA signed certificate, skip to step 3.
 - a Using SSH, log in to the SDDC Manager VM with the following credentials:

User name: `vcf`

Password: use the password specified in the deployment parameter workbook.
 - b Enter `su` to switch to the root user.
 - c Import the Redfish certificate from the Redfish VM to SDDC Manager VM by running the following command:


```
/opt/vmware/vcf/commonsvcs/scripts/cert-fetch-import-refresh.sh --ip=redfish-ip --port=port --service-restart=operationsmanager
```

`ip` Specify translation layer IP
`port` TLS/SSL port

The output displays information about the certificate to import including owner, issuer, serial number, validity, certificate fingerprints (md5, sha1, or sha256), signature algorithm name, subject, public key algorithm, and version. Verify this information.
 - d Answer the prompt.

Operations Manager is restarted. Wait for a few minutes and then proceed to step 4.

- 3 Restart Operations Manager:


```
systemctl restart operationsmanager
```

Wait for a few minutes before proceeding to the next step.
- 4 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
- 5 Enter the URL for the Redfish translation layer.
- 6 Enter the user name and password for the Redfish translation layer.
- 7 Click **Connect**.

Compose a Server

You can compose one or more servers by selecting the compute, network, and storage resources.

Prerequisites

- The translation layer must have been configured.
- The composed server must meet the minimum hardware requirements. See the *Planning and Preparation Workbook*.

Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
- 2 From the Available Resources table, select the zone where you want to compose a server. A zone corresponds to a physical boundary.
- 3 Click **Compose**.
- 4 In the Allocate Resources dialog box, select the compute for the server.
- 5 Select the storage.
- 6 Select the network interface.

The Choose number of servers section displays the number of servers you can compose based on the selected resources.
- 7 Select the number of servers you want to compose.
- 8 Click **Next**.
- 9 On the Review page, review the allocated resources to the servers.

Click **Back** to make any changes.
- 10 Click **Compose**.

Results

The compose server task is displayed in the Tasks table at the bottom of the Composable Infrastructure page. Click the name of the task for more information. When the server is composed, it is added to the Server Composition Summary table.

If Cloud Foundation does not receive a response from Redfish due to an external error, an error message is displayed. The hardware resources in the compose request are locked and cannot be used. You must free up these resources using the vendor UI. For more information, refer to the vendor documentation.

What to do next

- 1 Image the composed servers. See [Chapter 6 Installing ESXi Software on Cloud Foundation Servers](#).

- 2 Commission the composed servers. See [Commission Hosts](#).

View Composability Information

The Composable Infrastructure page displays information about available resources and composed servers.

Procedure

- ◆ On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
The Composable Infrastructure page appears. The Redfish translation layer information is displayed on the top of the page.
The Available Resources table displays the available zones and computer, storage, and network information available in each zone.
The Server Composition Summary table displays the composed servers.
The task panel at the bottom of the page shows the tasks performed and their status.

Add Storage

You can add storage to Dell MX composed servers.

Prerequisites

The server you are adding storage to must be unassigned.

Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
- 2 In the Composed Servers section, click the three dots next to the server you want to add storage to and click **Add Storage**.
The Allocate Storage dialog box displays the composed server details. Hovering over the Information icon next to the Current Storage field displays the SSD and HDD details of the server.
- 3 In the Add to Current Storage section, select the storage units you want to add and click **Next**.
- 4 On the Review page, verify the storage you are adding and click **Add**.

Remove Storage

You can remove storage units from a Dell MX composed server.

Prerequisites

The server you are removing storage from must be unassigned.

Procedure

- 1 On the Composable Infrastructure page, click the three dots next to the server you want to remove storage from.
- 2 Click **Remove Storage**.
- 3 For the storage type you want to remove, set the Remove from Allocated value to the appropriate value.
- 4 Click **Next**.
- 5 On the Review page, verify the storage you are removing and click **Remove**.

Decompose a Server

You can decompose a server that has not been assigned to a VI workload domain.

Prerequisites

The server to be decomposed must be unassigned.

Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
- 2 From the Server Composition Summary table, select the server to be decomposed.
- 3 Click **Decompose**.
- 4 In the Decompose Servers dialog box, click **Decompose**.

Monitoring Capabilities in the Cloud Foundation System

17

The Cloud Foundation system provides built-in capabilities to help you perform effective operations monitoring, troubleshooting, performance management, infrastructure capacity planning, and compliance monitoring and auditing.

You use the built-in monitoring capabilities for these typical scenarios.

Scenario	Examples
Are the systems online?	A host or other component shows a failed or unhealthy status.
Why did a storage drive fail?	Hardware-centric views spanning inventory, configuration, usage, and event history to provide for diagnosis and resolution.
Is the infrastructure meeting tenant service level agreements (SLAs)?	Analysis of system and device-level metrics to identify causes and resolutions.
At what future time will the systems get overloaded?	Trend analysis of detailed system and device-level metrics, with summarized periodic reporting.
What person performed which action and when?	History of secured user actions, with periodic reporting. Workflow task history of actions performed in the system.

The monitoring capabilities involve these features:

Tasks and subtasks

A task is a unit of work or a series of subtasks that perform an overall goal, such as creating a workload domain.

vRealize Log Insight instance deployed by Cloud Foundation

Use of the vRealize Log Insight instance deployed by Cloud Foundation is licensed separately. When this deployed vRealize Log Insight instance is licensed for use in your environment, and enabled in the SDDC Manager Dashboard, log content for the physical resources and the VMware SDDC virtual infrastructure are sent to the vRealize Log Insight instance. As a result, when you log in to the vRealize Log Insight Web interface, you can obtain a unified view of event and syslog information to assist with troubleshooting. Data from the events and audit events raised by Cloud Foundation is also sent to vRealize Log Insight. You can use the searching, query, and reporting features of vRealize Log Insight to create trend reports and

auditing reports from the event history. See [Using vRealize Log Insight Capabilities in Your Cloud Foundation System](#).

This chapter includes the following topics:

- [Viewing Tasks and Task Details](#)
- [Using vRealize Log Insight Capabilities in Your Cloud Foundation System](#)

Viewing Tasks and Task Details

From the SDDC Manager Dashboard, you can access all tasks. By default, the Dashboard displays the Recent Tasks widget, providing general information at a glance about the most recent tasks. A task is a unit of work or a series of subtasks that perform an overall goal, such as creating a workload domain.

In addition to the most recent tasks, you can view and search for all tasks by clicking **View All Tasks** at the bottom of the Recent Tasks widget. This opens the Tasks panel.

Note For more information about controlling the widgets that appear on the Dashboard page of the SDDC Manager Dashboard, see [Tour of the SDDC Manager User Interface](#).

Viewing and Filtering Task Details

The Tasks panel provides a high level view all tasks, displaying the descriptive task name, task status (for example, running, succeeded, or failed), and the timestamp for the last change in task status. You can also filter and search the task information as follows:

- Search tasks by clicking the filter icon in the Task column header and entering a search string.
- Filter tasks by status by clicking the filter icon in Status column. Select by category **All**, **Failed**, **Successful**, **Running**, or **Pending**.

Note Each category also displays the number of tasks with that status.

- Clear all filters by clicking **Reset Filter** at the top of the Tasks panel.
- Click **Refresh** to refresh the task list.

Note You can also sort the table by the contents of the Status and Last Occurrence columns.

Managing Tasks and Subtask Details

Expand a task to view details including the subtasks that comprise the task and their individual statuses.

- If a task is in a Failed state, you can also attempt to restart it by clicking **Restart Task**.

Note Not all tasks are restartable.

- If a task is in a Failed state, click on the icon next to the Failed status to view a detailed report on the cause.
- To view subtasks and their details, click **View Subtasks**.

Note You can filter subtasks in the same way you filter tasks.

Note You can also sort the table by the contents of the Status and Last Occurrence columns.

Resizing the Task Panel

Use the icons on the task panel to increase or decrease the panel size, or to close or reopen it.

Using vRealize Log Insight Capabilities in Your Cloud Foundation System

The vRealize Log Insight instance that is deployed by SDDC Manager is licensed separately. When the vRealize Log Insight instance is licensed for use in your Cloud Foundation environment, you can use the capabilities of vRealize Log Insight to work with the event and log data that is collected from the various hardware devices and SDDC virtual infrastructure.

vRealize Log Insight is a log aggregator that provides simplified log viewing and analysis. The vRealize Log Insight instance collects and indexes log content for the environment's physical resources and virtual infrastructure, and provides unified querying and analysis of the log content for problem diagnosis and repair. Similarly, SDDC Manager is configured by default to send all logs to vRealize Log Insight, enabling users to browse and search logs to troubleshoot SDDC Manager failures.

You can configure the vRealize Log Insight instance for remote syslog forwarding to an instance of vRealize Log Insight that is external to the Cloud Foundation system or to another syslog server. To configure vRealize Log Insight to forward events to a syslog target, see [Add vRealize Log Insight Event Forwarding Destination](#) in the vRealize Log Insight documentation.

To log in to the vRealize Log Insight Web interface from the SDDC Manager Dashboard, see [Enable vRealize Log Insight in Cloud Foundation](#).

Content Packs

The vRealize Log Insight instance includes a set of content packs. Content packs are read-only plug-ins to vRealize Log Insight that provide pre-defined knowledge about specific types of events such as log messages. The purpose of a content pack is to provide knowledge about a specific set of events in a format that is easily understandable by administrators, monitoring teams, and executives. A content pack consists of information that can be saved from either the Dashboards or Interactive Analytics pages in the vRealize Log Insight Web interface. Such information typically includes:

- Queries
- Fields

- Aggregations
- Alerts
- Dashboards

The vRealize Log Insight instance includes a number of VMware content packs, including the Cloud Foundation content pack. In the vRealize Log Insight web interface, these content packs display as widgets in the **Dashboards > VMware-VCF** page.

Content Pack	Overview
General	This content pack includes multiple subcategories of dashboards and analytics including overview, problems, event types, statistics, and agents.
VMware - NSX for vSphere	This content pack provides various dashboards and filters to give you insight into the data that is sent by the NSX for vSphere virtual infrastructure in the management and workload domains' vCenter Server instances.
VMware - Cloud Foundation	This content pack includes an overview dashboard that gives overall summary views of the data sent by the Cloud Foundation, and also provides detailed views for the various levels of interest, such as rack-level, server-level, switch-level, device-level, and so on.
VMware - vSAN	This content pack provides various dashboards and filters to give you insight into the logs that are sent by the management and workload domains' vSAN features.
VMware - vSphere	This content pack provides various dashboards and filters to give you insight into the data that is sent by the management and workload domains' vCenter Server instances.
VMware - vROPs	This content pack provides various dashboards and filters to give you insight into the logs that are sent by the management and workload domains' vRealize Operations features.

To see the dashboards for one of the content packs in the vRealize Log Insight Web interface, select **Dashboards** and then select the specific content pack dashboard in the left hand navigation bar.

Get Started Using the vRealize Log Insight Instance

Use of the vRealize Log Insight instance that is deployed by SDDC Manager is licensed separately. vRealize Log Insight delivers real-time log management for VMware environments, providing visibility of logs and easier troubleshooting across the physical and virtual infrastructure in your Cloud Foundation system.

During bring-up, SDDC Manager deploys and configures the vRealize Log Insight virtual appliance. From your deployed vRealize Log Insight instance, you can view and analyze logs to assist in troubleshooting, trend analysis, and so on.

The bring-up process also installs and configures content packs in the vRealize Log Insight instance. A content pack provides dashboards, extracted fields, predefined queries, and alerts that are related to the content pack's specific product or set of logs. When you launch the vRealize Log Insight Web interface, the installed content packs are ready for use. For an

overview of these content packs, see [Using vRealize Log Insight Capabilities in Your Cloud Foundation System](#). For detailed information on how to use the dashboards, predefined queries, and collected log data in vRealize Log Insight, see the [vRealize Log Insight product documentation](#).

You can open the vRealize Log Insight interface directly from the SDDC Manager Dashboard. For details, see [Enable vRealize Log Insight in Cloud Foundation](#).

If this is the first time after the initial bring-up process that the vRealize Log Insight Web interface is launched, type the system-assigned credentials into the login screen and then click **Login**. Then use the vRealize Log Insight Web interface to assign permissions to your superuser account and other user accounts.

Note You can look up the system-assigned credentials for the vRealize Log Insight Web interface by logging in to the SDDC Manager VM and running the `/home/vrack/bin/lookup-password` command.

Important Do not change the password of the admin account from within the vRealize Log Insight Web interface, or unpredictable results can occur. To change the admin account's password without rotating all account passwords, see [Manually Update Passwords](#).

Procedure

- 1 Open the vRealize Log Insight Web interface.
- 2 If the vRealize Log Insight login screen appears, log in with the appropriate credentials.
 - If this is the first time logging in to vRealize Log Insight after the initial bring-up process, use the username **admin** and the randomized password that was set when the passwords were rotated at the end of the bring-up process.
 - If you are using an account that was set up for you in vRealize Log Insight, use those credentials to log in.

When you are logging in to the vRealize Log Insight Web interface with the **admin** account after updating passwords, you must use the randomized password that is set for that account by the rotation procedure. For details about passwords, see [Manually Update Passwords](#).

Results

The vRealize Log Insight web interface appears with the display filtered to the **Dashboards > VMware-VCF > Overview** page to show the various event widgets.

Updating Cloud Foundation DNS and NTP Servers

18

If you need to make changes to the DNS or NTP server information that you provided during Cloud Foundation bring-up, you can use the SDDC Manager VM to update the servers starting from Cloud Foundation 3.7.2.

When you initially deploy Cloud Foundation, you complete the deployments parameter sheet to provide the system with the information required for bring-up. This includes up to two DNS servers and up to two NTP servers. You can change this server information at a later date, using the SDDC Manager VM.

This chapter includes the following topics:

- [Update DNS Server Configuration](#)
- [Update NTP Server Configuration](#)

Update DNS Server Configuration

Use this procedure to update the DNS server information that you provided during Cloud Foundation bring-up.

Cloud Foundation uses DNS servers to provide name resolution for various components in the system. You must provide root DNS domain information. Optionally, you can provide subdomain information. When you update the DNS server configuration, Cloud Foundation updates the components in a specific order:

- Platform Services Controllers
- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX Managers
- NSX-T Managers
- vRealize Suite Lifecycle Manager
- vRealize Log Insight
- vRealize Operations

- vRealize Automation

If the update fails, SDDC Manager rolls back the DNS settings for the failed component. Fix the underlying issue and retry the update starting with the failed component.

Note There is no rollback for the vRealize components. Check the logs, resolve any issues, and retry the update.

Updating the DNS server configuration is a disruptive process and can take some time to complete, depending on the size of your environment. Schedule DNS updates at a time that minimizes the impact to the system users.

Prerequisites

- Ensure that both forward and reverse DNS resolution is functional for each component using the updated DNS server.
- All Cloud Foundation components must be in an Active state.

Procedure

- 1 Using SSH, log in to the SDDC Manager VM using the **vcf** user account.
- 2 Get the current DNS server configuration information.

```
curl localhost/inventory/system-info | json_pp
```

- 3 Validate the new DNS server configuration information.

```
curl -H 'Content-Type: application/json' -X POST http://localhost/system/configurations/dns-servers/validator -d '{"dnsServers":[{"server":"<dns-server-ip>","isPrimary":"true"}]}' | json_pp
```

Replace *<dns-server-ip>* with the IP address of the new DNS server. Specify true or false for *isPrimary*, depending on whether or not the new DNS server is the primary DNS server.

The validator verifies forward and reverse name resolution for Cloud Foundation components using the new DNS server.

- 4 Monitor the status of the validation task.

```
curl http://localhost/operationsmanager/system/configurations/dns-servers/validator/status | json_pp
```

- 5 Check the result of the validation task.

```
curl http://localhost/operationsmanager/system/configurations/dns-servers/validator/result | json_pp
```

If validation succeeds, you can proceed to change the DNS server configuration. If validation fails, correct any issues and try again.

6 Change the DNS server configuration information.

```
curl -H 'Content-Type: application/json' -X POST http://localhost/system/configurations/dns-servers -d '{"dnsServers":[{"server":"<dns-server-ip>","isPrimary":"true"}]}' | json_pp
```

Replace *<dns-server-ip>* with the IP address of the new DNS server. Specify `true` or `false` for `isPrimary`, depending on whether or not the new DNS server is the primary DNS server.

Note the *<id>* that gets returned.

7 Track the status of the DNS update.

```
curl http://localhost/operationsmanager/workflows/<id> | json_pp
```

Replace *<id>* with the ID from the previous step. Wait for the task to complete.

8 Verify that the DNS configuration was updated.

```
curl localhost/inventory/system-info | json_pp
```

Update NTP Server Configuration

Use this procedure to update the NTP server information that you provided during Cloud Foundation bring-up.

Cloud Foundation uses NTP servers to synchronize time between the various components in the system. You must have at least one NTP server. When you update the NTP server configuration, Cloud Foundation updates the components in a specific order:

- Platform Services Controllers
- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX Managers
- NSX-T Managers
- vRealize Suite Lifecycle Manager
- vRealize Log Insight
- vRealize Operations
- vRealize Automation

If the update fails, SDDC Manager rolls back the NTP settings for the failed component. Fix the underlying issue and retry the update starting with the failed component.

Note There is no rollback for the vRealize components. Check the logs, resolve any issues, and retry the update.

Updating the NTP server configuration is a disruptive process and can take some time to complete, depending on the size of your environment. Schedule NTP updates at a time that minimizes the impact to the system users.

Prerequisites

- Any new NTP server is reachable by all components.
- Time skew between new NTP servers is less than 5 minutes.

Procedure

- 1 Using SSH, log in to the SDDC Manager VM using the **vcf** user account.
- 2 Get the current NTP server configuration information.

```
curl localhost/inventory/system-info | json_pp
```

- 3 Validate the new NTP server configuration information.

```
curl -H 'Content-Type: application/json' -X POST http://localhost/system/configurations/ntp-servers/validator -d '{"ntpServers":[{"server":"<ntp-server-ip>"}]}' |json_pp
```

Replace *<ntp-server-ip>* with the IP address of the new NTP server. To enter multiple NTP servers, use the following format: `{"ntpServers":[{"server":"<ntp-server-ip-1>","server":"<ntp-server-ip-2>"}]}`.

The validator verifies that the Cloud Foundation components can communicate with the new NTP server.

- 4 Monitor the status of the validation task.

```
curl http://localhost/operationsmanager/system/configurations/ntp-servers/validator/status | json_pp
```

- 5 Check the result of the validation task.

```
curl http://localhost/operationsmanager/system/configurations/ntp-servers/validator/result | json_pp
```

If validation succeeds, you can proceed to change the NTP server configuration. If validation fails, correct any issues and try again.

- 6 Change the NTP server configuration information.

```
curl -H 'Content-Type: application/json' -X POST http://localhost/system/configurations/ntp-servers -d '{"ntpServers":[{"server":"<ntp-server-ip>"}]}' |json_pp
```

Replace *<ntp-server-ip>* with the IP address of the new NTP server. To enter multiple NTP servers, use the following format: `{"ntpServers":[{"server":"<ntp-server-ip-1>","server":"<ntp-server-ip-2>"}]}`.

Note the *<id>* that gets returned.

7 Track the status of the NTP update.

```
curl http://localhost/tasks/registrations/<id> | json_pp
```

Replace *<id>* with the ID from the previous step. Note the *<taskURL>*.

```
curl <taskURL> | json_pp
```

Wait for the task to complete.

8 Verify that the NTP configuration was updated.

```
curl localhost/inventory/system-info | json_pp
```

Supportability and Serviceability (SoS) Utility

19

The SoS utility is a command-line Python tool that can be used for the following:

- Run health checks.
- On-demand vSAN partition cleanup.
- Collect logs for Cloud Foundation components.
- On-demand host cleanup.

To run the SoS utility, SSH in to the SDDC Manager VM using the **vcf** administrative user account. For basic operations, enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-1 --option-2 --option-3 ... --option-n
```

To list the available command options, use the `--help` long option or the `-h` short option.

```
sudo /opt/vmware/sddc-support/sos --help
sudo /opt/vmware/sddc-support/sos -h
```

Note You can specify some options in the conventional GNU/POSIX syntax, using `--` for the long option and `-` for the short option.

For privileged operations, enter **su** to switch to the root user, and navigate to the `/opt/vmware/sddc-support` directory and type `./sos` followed by the options required for your desired operation.

This chapter includes the following topics:

- [SoS Utility Options](#)
- [Collect Logs for Your Cloud Foundation System](#)

SoS Utility Options

This section lists the specific options you can use with the SoS utility.

SoS Utility Help Options

Use these options to see information about the SoS utility itself. For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the **vcf** password when prompted.

Option	Description
--help	Provides a summary of the available SoS utility options
-h	
--version	Provides the SoS utility's version number.
-v	

SoS Utility VMware Cloud Foundation Summary Options

These options provide information about the Cloud Foundation system and tasks. For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the **vcf** password when prompted.

Option	Description
--get-vcf-summary	Returns information about your Cloud Foundation system, including CEIP, domains and clusters, hosts, licensing, network pools, SDDC Manager, VCF services, and solutions (vRealize Log Insight, vRealize Automation, and so on).
--get-vcf-tasks-summary	Returns information about Cloud Foundation tasks, including the time the task was created and the status of the task.

SoS Utility Generic Options

These are generic options for the SoS utility. For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the **vcf** password when prompted.

Note For generic options related to log collection, see [Collect Logs for Your Cloud Foundation System](#).

Option	Description
<code>--ceip-tagging-get</code>	Returns setting for the VMware CEIP program. For information about the program, see Chapter 3 Configuring Customer Experience Improvement Program .
<code>--ceip-tagging-set</code>	Enrolls your deployment in the CEIP program.
<code>--configure-sftp</code>	Configures SFTP for logs.
<code>--debug-mode</code>	Runs the SoS utility in debug mode.
<code>--domain-name DOMAINNAME</code>	Specify the name of the workload domain name on which the SoS operation is to be performed. To run the operation on all domains, specify <code>--domain-name ALL</code> .
	Note If you omit the <code>--domain-name</code> flag and domain name, the SoS operation is performed only on the management domain.
<code>--force</code>	Allows SoS operations to be formed while workflows are running. Note It is recommended that you do not use this option.
<code>--history</code>	Displays the last 20 SoS operations performed.
<code>--ondemand-service</code>	Include this flag to execute commands on all ESXi hosts in a domain. Warning Contact VMware support before using this option.
<code>--ondemand-service-json JSON file path</code>	Include this flag to execute commands in the JSON format on all ESXi hosts in a domain. For example, <code>/opt/vmware/sddc-support/<JSON file name></code>
<code>--setup-json SETUPJSON</code>	Custom setup-json file for log collection. SoS prepares the inventory automatically based on the environment where it is running. If you want to collect logs for a pre-defined set of components, you can create a <code>setup.json</code> file and pass the file as input to SoS. A sample JSON file is available on the SDDC Manager VM in the <code>/opt/vmware/sddc-support/</code> directory.
<code>--skip-known-host-check</code>	Skips the specified check for SSL thumbprint for host in the known host.
<code>--zip</code>	Creates a zipped TAR file for the output.

SoS Utility Options for Health Check

These SoS commands are used for checking the health status of various components or services, including connectivity, compute, storage, database, domains, and networks. For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the **vcf** password when prompted.

A green status indicates that the health is normal, yellow provides a warning that attention might be required, and red (critical) indicates that the component needs immediate attention.

Option	Description
<code>--json-output-dir JSONDIR</code>	Outputs the results of any health check as a JSON file to the specified directory, JSONDIR.
<code>--certificate-health</code>	Verifies that the component certificates are valid (within the expiry date).
<code>--connectivity-health</code>	Performs a connectivity health check to inspect whether the different components of the system such as the ESXi hosts, Virtual Center Servers, Inventory Service VMs, Log Insight VM, NSX Manager VMs, PSC VMs, SDDC Manager VM can be pinged.
<code>--composability-infra-health</code>	Performs an API connectivity health check of the composable infrastructure. If no composable infrastructure exists, this flag is ignored. If found, the utility checks connectivity status through the composable infrastructure API, such as Redfish.
<code>--compute-health</code>	Performs a compute health check.
<code>--general-health</code>	Verifies ESXi entries across all sources, checks the Postgres DB operational status for hosts, checks ESXi for error dumps, and gets NSX Manager and cluster status.
<code>--get-host-ips</code>	Returns server information.
<code>--get-inventory-info</code>	Returns in a tabular format inventory details for the specified Cloud Foundation component, such as Platform Services Controller/vCenter Server NSX, and ESXi. Optionally, add the flag <code>--domain name ALL</code> to return all details.
<code>--health-check</code>	Performs all available health checks.
<code>--ntp-health</code>	Verifies whether the time on the components is synchronized with the NTP server in the SDDC Manager VM. It also ensures that the hardware and software timestamp of ESXi hosts are within 5 minutes of the SDDC Manager VM.
<code>--password-health</code>	Returns the status of all current passwords, such as Last Changed Date, Expiry Date, and so on.
<code>--services-health</code>	Performs a services health check to confirm whether services within the Inventory Service VM and within SDDC Manager (like Lifecycle Management Server) are running.
<code>--storage-health</code>	Performs a check on the vSAN disk health of the ESXi hosts and vCenter clusters. Also runs proactive vSAN tests to verify the ability to create VMs within the vSAN disks.
<code>--run-vsan-checks</code>	Runs proactive vSAN tests to verify the ability to create VMs within the vSAN disks.

SoS Utility Options for vSAN Stretched Clusters

Use create a vSAN stretched cluster, convert a vSAN stretch cluster to a standard vSAN cluster, and add/replace hosts in a vSAN stretched cluster. See [Chapter 15 Stretching Clusters](#). For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account, enter **su** to switch to the root user, navigate to the `/opt/vmware/sddc-support` directory, and type the following command:

```
./sos --option-name
```

Option	Description
<code>--expand-stretch-cluster</code>	Add hosts to or replace a host in a vSAN stretch cluster. Used with <code>--sc-domain --sc-cluster --sc-hosts --esxi-license-key</code> . For example, <code>--expand-stretch-cluster --sc-domain MGMT --sc-cluster SDDC-Cluster1 --sc-hosts esxi-9.vrack.vsphere.local, esxi-10.vrack.vsphere.local --esxi-license-key AAAAA-BBBBB-CCCC-DDDD-EEEE</code> .
<code>--show-clusters</code>	Shows all domains and clusters.
<code>--show-free-hosts</code>	Shows all free hosts.
<code>--stretch-vsan</code>	Create a vSAN stretch cluster. Used with <code>--sc-domain --sc-cluster --sc-hosts --witness-host-fqdn --witness-vsan-ip --witness-vsan-cidr --esxi-license-key</code> . For example, <code>--stretch-vsan --sc-domain MGMT --sc-cluster SDDC-Cluster1 --sc-hosts esxi-5.vrack.vsphere.local, esxi-6.vrack.vsphere.local --witness-host-fqdn esxi-11.vrack.vsphere.local --witness-vsan-ip 10.0.12.96 --witness-vsan-cidr 10.0.12.0/24 --esxi-license-key AAAAA-BBBBB-CCCC-DDDD-EEEE</code> , where AAAAA-BBBBB-CCCC-DDDD-EEEE is a valid ESXi license key.
<code>--sc-domain SCDOMAIN</code>	Specify the domain, SCDOMAIN, to use for stretched vSAN. For example, <code>--sc-domain MGMT</code> .
<code>--sc-cluster SCCLUSTER</code>	Specify the cluster, SCCLUSTER, to use for stretched vSAN. For example, <code>--sc-cluster SDDC-Cluster1</code> .
<code>--sc-hosts SCHOSTS [SCHOST1, SCHOST2 ...]</code>	Specify the hosts, SCHOSTS, to use for stretched vSAN. For example, <code>--sc-hosts esxi-5.vrack.vsphere.local, esxi-6.vrack.vsphere.local</code> .
<code>--witness-host-fqdn WITNESSHOSTFQDN</code>	Specify the fully qualified domain name, WITNESSHOSTFQDN, of the witness host. For example, <code>--witness-host-fqdn esxi-11.vrack.vsphere.local</code> .
<code>--witness-vsan-ip WITNESSHOSTVSANIP</code>	Specify the IP address, WITNESSHOSTVSANIP, of the witness host. For example, <code>--witness-vsan-ip 10.0.12.96</code> .
<code>--witness-vsan-cidr WITNESSHOSTVSANCIDR</code>	Specify the Classless Inter-Domain Routing (CIDR) block, WITNESSHOSTVSANCIDR, of the witness host. For example, <code>--witness-vsan-cidr 10.0.12.0/24</code> .
<code>--esxi-license-key ESXILICENSEKEY</code>	Specify the license key, ESXILICENSEKEY, to use for ESXi hosts. For example, <code>--esxi-license-key AAAAA-BBBBB-CCCC-DDDD-EEEE</code> .

SoS Utility Options for Fixing vSAN Partitions

Use this option to clean up vSAN partitions on one or more ESXi hosts. These options can be run only from the SDDC Manager VM. For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account, enter **su** to switch to the root user, navigate to the `/opt/vmware/sddc-support` directory, and type the following command:

```
./sos --option-name
```

Option	Description
<code>--cleanup-vsan</code>	Cleans up vSAN Partitions in ESXi hosts. Optionally, you can specify the ESXi hosts, by IP address, to run the vSAN cleanup. Use commas (with no spaces) to separate multiple IP addresses.

SoS Utility Options for Managing ESXi Hosts

Use these options to clean up and manage ESXi hosts, including enabling SSH, cleaning up dirty hosts, and locking down hosts. For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account, enter **su** to switch to the root user, navigate to the `/opt/vmware/sddc-support` directory, and type the following command:

```
./sos --option-name
```

Option	Description
<code>--cleanup-decommissioned-host</code>	Performs clean-up on the specified, decommissioned ESXi hosts by passing the JSON. For example: <code>--cleanup-decommissioned-host /opt/vmware/sddc-support/decommissioned_host_cleanup_sample.json</code>
<code>--cleanup-host</code>	<p>Performs clean-up on all or specified dirty ESXi hosts.</p> <ul style="list-style-type: none"> ■ To clean up all dirty hosts, include ALL: <code>--cleanup-host ALL</code>. ■ To specify multiple dirty hosts, separate the IP addresses with a comma: <code>--cleanup-host 10.0.0.4,10.0.0.5,10.0.0.6</code>. <p>Note A dirty host is a host that has been removed from a cluster in a workload domain. A dirty host cannot be assigned to another workload domain until it is cleaned up.</p>
<code>--disable-lockdown-esxi</code>	<p>Disables lockdown mode on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> ■ To disable lockdown on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>. ■ To disable lockdown on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>. <p>Note If you do not specify domain, this command affects only the MGMT domain by default.</p>
<code>--enable-lockdown-esxi</code>	<p>Enables lockdown mode on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> ■ To enable lockdown on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>. ■ To enable lockdown on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>. <p>Note If you do not specify domain, this command affects only the MGMT domain by default.</p>

Option	Description
<code>--disable-ssh-esxi</code>	<p>Disables SSH on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> ■ To disable SSH on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>. ■ To disable SSH on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>. <p>Note If you do not specify domain, this command affects only the MGMT domain by default.</p>
<code>--enable-ssh-esxi</code>	<p>Enables SSH on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> ■ To enable SSH on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>. ■ To enable SSH on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>. <p>Note If you do not specify domain, this command affects only the MGMT domain by default.</p>

SoS Utility Options for vRealize Suite Lifecycle Manager

Use these options to redeploy vRealize Suite Lifecycle Manager and monitor the redeployment. For these options, SSH in to the SDDC Manager VM using the **vcf** administrative user account, enter **su** to switch to the root user, navigate to the `/opt/vmware/sddc-support` directory, and type the following command:

```
./sos --option-name
```

Note You should only redeploy vRealize Suite Lifecycle Manager when directed to do so by VMware Support.

Option	Description
<code>--vrs1cm-redeploy</code>	Redeploys vRealize Suite Lifecycle Manager. Provides a taskID for the operation.
<code>--get-vrs1cm-redeploy-task-status <taskID></code>	Returns vRealize Suite Lifecycle Manager redeployment status for the specified taskID.

Collect Logs for Your Cloud Foundation System

Use the SoS utility to collect the logs for various software components in the system.

Use these options when retrieving support logs from your environment's various components.

- To collect all logs from all components, you can run the SoS utility without specifying any component-specific options.
- To collect logs for a specific component, run the utility with the appropriate options.

For example, the `--domain-name` option is important. If omitted, the SoS operation is performed only on the management domain. See [SoS Utility Options](#).

Log files for the vRealize Log Insight agent in vCenter Server are collected when vCenter Server log files are collected.

After running the SoS utility, you can examine the resulting logs to troubleshoot issues, or provide to VMware Technical Support if requested. VMware Technical Support might request these logs to help resolve technical issues when you have submitted a support request. The diagnostic information collected using the SoS utility includes logs for the various VMware software components and software products deployed in your Cloud Foundation environment.

Table 19-1. SoS Utility Log File Options

Option	Description
<code>--api-logs</code>	Collects output from REST endpoints for SDDC Manager inventory and LCM.
<code>--cassandra-logs</code>	Collects logs from the Apache Cassandra database only. <code>cassandra-bundle.tgz</code> contains Cassandra nodetool and debug logs. Apache Cassandra processes run in each of the infrastructure virtual machines.
<code>--dump-only-sddc-java-threads</code>	Collects only the Java thread information from the SDDC Manager.
<code>--esx-logs</code>	Collects logs from the ESXi hosts only. Logs are collected from each ESXi host available in the deployment.
<code>--li-logs</code>	Collects logs from vRealize Log Insight VMs only.
<code>--log-dir LOGDIR</code>	Specifies the directory to store the logs.
<code>--log-folder LOGFOLDER</code>	Specifies the name of the log directory.
<code>--no-clean-old-logs</code>	Use this option to prevent the utility from removing any output from a previous collection run. By default, the SoS utility. By default, before writing the output to the directory, the utility deletes the prior run's output files that might be present. If you want to retain the older output files, specify this option.
<code>--no-health-check</code>	Skips the health check executed as part of log collection.
<code>--nsx-logs</code>	Collects logs from the NSX Manager, NSX Controller, and NSX Edge instances only.
<code>--psc-logs</code>	Collects logs from the Platform Services Controller instances only.
<code>--rvc-logs</code>	Collects logs from the Ruby vSphere Console (RVC) only. RVC is an interface for ESXi and vCenter. Note If the Bash shell is not enabled in vCenter, RVC log collection will be skipped . Note RVC logs are not collected by default with <code>./sos</code> log collection. You must enable RVC to collect RVC logs.
<code>--sddc-manager-logs</code>	Collects logs from the SDDC Manager only. <code>sddc<timestamp>.tgz</code> contains logs from the SDDC Manager file system's etc, tmp, usr, and var partitions.
<code>--test</code>	Collects test logs by verifying the files.
<code>--vc-logs</code>	Collects logs from the vCenter Server instances only. Logs are collected from each vCenter server available in the deployment.
<code>--vdi-logs</code>	Collects logs for Horizon domain management components .
<code>--vrealize-logs</code>	Collects logs from vRealize components deployed in the system (vRealize Suite Lifecycle Manager, vRealize Log Insight, vRealize Operations, and vRealize Automation).

Procedure

- 1 Using SSH, log in to the SDDC Manager VM with the following credentials:

Username: **vcf**

Password: use the password specified in the deployment parameter sheet

- 2 To collect the logs, run the SoS utility without specifying any component-specific options.

```
sudo /opt/vmware/sddc-support/sos
```

Enter the **vcf** password when prompted.

To collect logs for a specific component, run the utility with the appropriate options.

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Note By default, before writing the output to the directory, the utility deletes the prior run's output files that might be present. If you want to retain the older output files, specify the `--no-clean-old-logs` option.

If you do not specify the `--log-dir` option, the utility writes the output to the `/var/log/vmware/vcf/sddc-support` directory in the SDDC Manager VM

Results

The utility collects the log files from the various software components in all of the racks and writes the output to the directory named in the `--log-dir` option. Inside that directory, the utility generates output in a specific directory structure.

Example

```
vcf@sddc-manager [ ~ ]$ sudo /opt/vmware/sddc-support/sos --domain-name MGMT --skip-known-host-check
--log-dir /tmp/new
[sudo] password for vcf
Welcome to Supportability and Serviceability(SoS) utility!
Performing SoS operation for MGMT domain components
Logs : /tmp/new/sos-2019-09-03-21-04-40-11793
Log file : /tmp/new/sos-2019-09-03-21-04-40-11793/sos.log
Log Collection completed successfully for : [HEALTH-CHECK, SDDC-MANAGER, NSX-MANAGER, PSC, API-LOGS,
ESX, LOGINSIGHT, VMS-SCREENSHOT, VCENTER-SERVER, VCF-SUMMARY]
```

What to do next

Change to the output directory to examine the collected log files.

Component Log Files Collected By the SoS Utility

The SoS utility writes the component log files into an output directory structure within the file system of the SDDC Manager instance in which the command is initiated, for example:

```
vcf@sddc-manager [ ~ ]$ sudo /opt/vmware/sddc-support/sos
[sudo] password for vcf
Welcome to Supportability and Serviceability(SoS) utility!
Performing SoS operation for MGMT domain components
Logs : /var/log/vmware/vcf/sddc-support/sos-2019-09-03-20-55-41-10053
Log file : /var/log/vmware/vcf/sddc-support/sos-2019-09-03-20-55-41-10053/sos.log
NOTE : The Health check operation was invoked without --skip-known-host-check, and so will skip Connectivity Health,
Password Health and Certificate Health Checks because of security reasons.

Log Collection completed successfully for : [HEALTH-CHECK, SDDC-MANAGER, NSX_MANAGER, PSC, API-LOGS, ESX,
LOGINSIGHT, VMS-SCREENSHOT, VCENTER-SERVER, VCF-SUMMARY]
```

esx Directory Contents

In each rack-specific directory, the esx directory contains the following diagnostic files collected for each ESXi host in the rack:

File	Description
<code>esx-IP-address.tgz</code>	Diagnostic information from running the <code>vm-support</code> command on the ESXi host. An example file is <code>esx-192.168.100.101.tgz</code> .
<code>SmartInfo-IP-address.txt</code>	S.M.A.R.T. status of the ESXi host's hard drive (Self-Monitoring, Analysis, and Reporting Technology). An example file is <code>SmartInfo-192.168.100.101.txt</code> .
<code>vsan-health-IP-address.txt</code>	vSAN cluster health information from running the standard command <code>python /usr/lib/vmware/vsan/bin/vsan-health-status.py</code> on the ESXi host. An example file is <code>vsan-health-192.168.100.101.txt</code> .

loginsight Directory Contents

The loginsight directory contains diagnostic information files collected from the vRealize Log Insight cluster. The support bundle for each node is collected from the cluster's load balancer VM.

File	Description
<code>load-balancer.vrack.vsphere.local-loginsight-support.tgz</code>	Compressed TAR file consisting of support bundles collected from each node in the vRealize Log Insight cluster. For example: <code>loginsight-loginsight-node-<i><node-number></i>.vrack.vsphere.local-<i><time-stamp></i></code> .
<code>loginsight-loginsight-node-<i><node-number></i>.vrack.vsphere.local-<i><time-stamp></i></code>	Contains the following: <code>README</code> , <code>boot</code> , <code>error.log</code> , etc, <code>proc</code> , <code>usr</code> , <code>action.log</code> , <code>commands</code> , <code>errors-ignored.log</code> , <code>opt</code> , <code>storage</code> , and <code>var</code> .

nsx Directory Contents

In each rack-specific directory, the nsx directory contains the diagnostic information files collected for the NSX Manager, NSX Controller, and NSX Edge instances deployed in that rack.

The number of files in this directory depends on the number of NSX Manager, NSX Controller, and NSX Edge instances that are deployed in the rack. In a given rack, each management domain has one NSX Manager instance and a minimum of three NSX Controller instances, and any VI workload domains in the rack each have one NSX Manager instance and at least three NSX Controller instances. NSX Edge instances are only deployed to support vRealize Operations and vRealize Automation, which are optional components.

File	Description
VMware-NSX-Manager-tech-support- <i>nsxmanagerIPAddr</i> .tar.gz	Standard NSX Manager compressed support bundle, generated using the NSX for vSphere API POST https://nsxmanagerIPAddr/api/1.0/appliance-management/techsupportlogs/NSX , where <i>nsxmanagerIPAddr</i> is the IP address of the NSX Manager instance. An example is VMware-NSX-Manager-tech-support-10.0.0.8.tar.gz.
VMware-NSX-Controller-tech-support- <i>nsxmanagerIPAddr</i> - <i>controller-controllerId</i> .tgz	Standard NSX Controller compressed support bundle, generated using the NSX for vSphere API to query the NSX Controller technical support logs: GET https://nsxmanagerIPAddr/api/2.0/vdn/controller/controllerId/techsupportlogs , where <i>nsxmanagerIPAddr</i> is the IP address of the NSX Manager instance and <i>controllerId</i> identifies the NSX Controller instance. Examples are VMware-NSX-Controller-tech-support-10.0.0.8-controller-1.tgz, VMware-NSX-Controller-tech-support-10.0.0.8-controller-2.tgz, VMware-NSX-Controller-tech-support-10.0.0.8-controller-3.tgz.
VMware-NSX-Edge-tech-support- <i>nsxmanagerIPAddr</i> - <i>edgeId</i> .tgz	Standard NSX Edge support bundle, generated using the NSX for vSphere API to query the NSX Edge support logs: GET https://nsxmanagerIPAddr/api/4.0/edges/edgeId/techsupportlogs , where <i>nsxmanagerIPAddr</i> is the IP address of the NSX Manager instance and <i>edgeId</i> identifies the NSX Edge instance. An example is VMware-NSX-Edge-tech-support-10.0.0.7-edge-1.log.gz.
Note This information will only be collected if NSX Edges are deployed.	

psc Directory Contents

In the rack-1 directory, the psc directory contains the diagnostic information files collected for the Platform Services Controller instances deployed in that rack.

File	Description
vm-support- <i>psclIPAddr</i> .tar.gz	Standard Platform Services Controller support bundle downloaded from the Platform Services Controller instance with IP address <i>psclIPAddr</i> .

vc Directory Contents

In each rack-specific directory, the vc directory contains the diagnostic information files collected for the vCenter Server instances deployed in that rack.

The number of files in this directory depends on the number of vCenter Server instances that are deployed in the rack. In a given rack, each management domain has one vCenter Server instance, and any VI workload domains in the rack each have one vCenter Server instance.

File	Description
vc- <i>vcsaFQDN</i> - <i>timestamp</i> .tgz	Standard vCenter Server support bundle downloaded from the vCenter Server Appliance instance having a fully-qualified domain name <i>vcsaFQDN</i> . The support bundle is obtained from the instance using the standard <code>vc-support.sh</code> command.

Managing Shutdown and Startup of Cloud Foundation

20

You might have situations in which you want to shut down and start up the system. In such situations, you must start up and shut down the management virtual machines according to a predefined order.

The following situations require shutting down and starting up the Cloud Foundation system:

- Performing patch or upgrade operations of SDDC Manager applications.
- Performing recovery or failover operations of SDDC Manager applications.
- Performing imaging at one location and shipping the rack for deployment at another location.

This chapter includes the following topics:

- [Shut Down a Cloud Foundation System](#)
- [Start Up a Cloud Foundation System](#)

Shut Down a Cloud Foundation System

You must shut down the system components in a strict order to avoid data loss and faults in the components.

Prerequisites

- Verify that you have direct console access to the switches and ESXi hosts in the system.
- Coordinate the shutdown in advance with business stakeholders to minimize any impact.
- Verify that no VMs are running on snapshots.
- Verify that you have saved the account passwords to a location external from the Cloud Foundation system you are shutting down. See [Look Up Account Credentials](#) .
- Verify that valid backups of all management VMs, tenant VMs, and switch configurations are available and saved to a location external from the Cloud Foundation system you are shutting down.
- If a data protection solution is running on any of the domains, verify that it is properly shut down according to the vendor instructions.

- See Knowledge Base article 2142676 [Shutting down and powering on a vSAN 6.x Cluster when vCenter Server is running on top of vSAN](#) for information about verifying the state of the vSAN cluster before a shutdown.

Procedure

- 1 Before starting the shutdown procedure, note down the following information:
 - The hostname and IP address of the ESXi hosts that are members of the management domain. To see the hosts in the management domain, navigate to the **Hosts** tab on the **Domain Details** page in the SDDC Manager Dashboard.
 - The hostname and IP address of the ESXi hosts that are members of each workload domain. To see the hosts in the VI workload domains, for each domain navigate to the **Hosts** tab on the **Domain Details** page in the SDDC Manager Dashboard.
- 2 Shut down the VMs in each Horizon domain (if any).
 - a On the SDDC Manager Dashboard, navigate to the Horizon domain.
 - b Click the launch link () for the vCenter Server instance that is displayed in the **Service** tabs in the domain details window for that Horizon domain.

A new browser tab opens and displays the vSphere Web Client.
 - c Locate the VMs for that Horizon domain.
 - d Shut down these VMs.
 - e Perform the above steps for each Horizon domain.
- 3 Shut down the workload VMs in each VI workload domain.
 - a On the SDDC Manager Dashboard, navigate to the workload domain.
 - b Click the launch link () for the vCenter Server instance that is displayed in the **Service** tabs in the domain details window for that workload domain.

A new browser tab opens and displays the vSphere Web Client.
 - c Locate the VMs for that workload domain.
 - d Shut down these VMs.

Note Each workload domain includes a three-node NSX Controller cluster. Shut down these VMs last.

- e Perform the above steps for each VI workload domain.

4 Place the hosts for each workload domain in maintenance mode.

You must use the ESXCLI command, which supports setting the vSAN mode when entering maintenance mode.

- a For each ESXi host, connect and log in to the ESXi console using SSH.
- b Place each host into maintenance mode using the following command, with the `noAction` option included.

```
esxcli system maintenanceMode set -e true -m noAction
```

- c After a few minutes, confirm each host is in maintenance mode by repeating the command.

```
esxcli system maintenanceMode set -e true -m noAction
```

It should return the following:

```
Maintenance mode is already enabled.
```

- d Shut down all the ESXi hosts in the VI workload domain.

```
# poweroff
```

- e Repeat the above steps for each VI workload domain.

5 Shut down the vRealize Suite Lifecycle Manager appliance in the management domain.

6 Shut down the vRealize Log Insight virtual appliances in the management domain in the following order:

Important Verify that the console of each virtual appliance and its services are fully shut down before proceeding to the next step.

- All vRealize Log Insight Worker nodes.
- The vRealize Log Insight Master node.

7 Shut down the vRealize Operations Manager virtual appliances in the management domain in the following order:

Important Verify that the console of each virtual appliance and its services are fully shut down before proceeding to the next step.

- All vRealize Operations Manager Remote Collector nodes.
- All vRealize Operations Manager Data nodes.
- The vRealize Operations Manager Replica node.
- The vRealize Operations Manager Master node.

- 8 Shut down the vRealize Automation virtual appliance and IaaS components in the management domain in the following order:

Important Verify that the console of each virtual appliance or VM and its services are fully shut down before proceeding to the next step.

- All vRealize Automation IaaS Distributed Execution Management (DEM) VMs.
- All vRealize Automation IaaS Proxy Agent VMs.
- All vRealize Automation IaaS Manager Server VMs.

Note Shut down the secondary IaaS Manager Server VM first; shut down the primary IaaS Manager Server VM second.

- All vRealize Automation IaaS Web Server VMs.

Note Shut down the secondary IaaS Web Server VM first; shut down the primary IaaS Web Server VM second.

- All vRealize Automation virtual appliances.
- The vRealize Automation IaaS SQL Server VM.

- 9 Shut down the infrastructure management virtual appliances in the management domain in the following sequence.

- a Shut down the following virtual appliances using the SSH console, verifying that the console of each virtual appliance and its services are fully shut down before proceeding to the next step.

- All NSX Edge Service Gateway virtual appliances.
- The NSX Manager virtual appliances for the VI workload domains.
- The NSX Manager virtual appliance for the management domain.
- All NSX Controller cluster virtual appliances for the management domain.

- b Shut down the remaining virtual appliances or VMs from their hosts in the ESXi Host Client on each management ESXi host.

- The vCenter Server virtual appliance for the management domain.
- The SDDC Manager VM.
- The Platform Services Controller virtual appliances.

10 Place the management domain ESXi hosts in maintenance mode.

You must use the ESXCLI command that supports setting the vSAN mode when entering maintenance mode.

- a For each ESXi host, connect and log in to the ESXi console using SSH.
- b Put each host into maintenance mode using the following command, with the `noAction` option included.

```
# esxcli system maintenanceMode set -e true -m noAction
```

- c After a few minutes, confirm each host is in maintenance mode by repeating the command.

```
esxcli system maintenanceMode set -e true -m noAction
```

It should return the following:

```
Maintenance mode is already enabled.
```

- d Shut down all the ESXi hosts in the management domain.

```
# poweroff
```

11 Shut down the unassigned ESXi hosts in the Cloud Foundation system, if any.

- a For each unassigned ESXi host, connect and log in to the ESXi console using SSH.
- b Shut down each unassigned ESXi host.

```
# poweroff
```

Start Up a Cloud Foundation System

You must start up the system components of the system in a strict order to avoid data loss and faults in the components.

Prerequisites

- Verify that you have direct console access to the switches and ESXi hosts in the system.
- Verify that you have the host names and IP addresses of the ESXi hosts that are members of the management domain.

You can obtain this information in the Domain Details pages in the SDDC Manager Dashboard.

- Verify that you have the host names and IP addresses of the ESXi hosts that are members of each VI workload domain.

You can obtain this information in the Domain Details pages in the SDDC Manager Dashboard.

- See Knowledge Base article 2142676 [Shutting down and powering on a vSAN 6.x Cluster when vCenter Server is running on top of vSAN](#) for information about starting up hosts and exiting maintenance mode.

Procedure

- 1 Power on each ESXi host in the management domain, and exit maintenance mode.

- a Use SSH to connect and log in to the ESXi console.
- b Use the following CLI command to exit maintenance mode.

```
# esxcli system maintenanceMode set -e false
```

- c Perform the above steps on each ESXi host until none are in maintenance mode.

- 2 Power on each ESXi host in the first VI workload domain, and exit maintenance mode.

- a Use SSH to connect and log in to the ESXi console.
- b Use the following CLI command to exit maintenance mode.

```
# esxcli system maintenanceMode set -e false
```

- c Perform the above steps on each ESXi host until none are in maintenance mode.

- d Repeat the above steps for each VI workload domain.

- 3 Power on the infrastructure management VMs in the management domain.

Important You must power on the VMs using the ESXi host client on each management ESXi host.

Important You must wait until each VM is powered on and all its services started before powering on the next VM.

Power on the VMs in the following order:

- Platform Services Controller virtual appliances.
- vCenter Server for the management domain.
- SDDC Manager VM.
- NSX Manager virtual appliance for the management domain.
- NSX Controller cluster virtual appliances for the management domain.
- NSX Edge Service Gateway virtual appliances.
- vCenter Server for each VI workload domain.
- NSX Manager virtual appliance for each VI workload domain.

- 4 Log in to the SDDC Manager Dashboard to verify that it displays correctly.
 - a On the SDDC Manager Dashboard, navigate to the management domain.
 - b In the domain details panel, click the launch for the vCenter Server instance.

A new browser tab opens and displays the vSphere Web Client.
- 5 Power on the vRealize Automation virtual appliance and IaaS components in the management domain.

Important You must wait until each VM or virtual appliance is powered on and all of its services started before running on the next VM.

Power on the VMs in the following order:

- The vRealize Automation IaaS SQL Server VM.
- All vRealize Automation virtual appliances.
- All vRealize Automation IaaS Web Server VMs.

Note Power on the primary IaaS Web Server VM first.

- All vRealize Automation IaaS Manager Services.

Note Power on the primary IaaS Manager Server VM first.

- All vRealize Automation IaaS proxy agents.
- All vRealize Automation IaaS Distributed Execution Management (DEM) hosts.

- 6 Power on the vRealize Operations Manager virtual appliances in the management domain.

Important You must wait until each VM or virtual appliance is powered on and all of its services running before powering on the next VM.

Power on the VMs in the following order:

- The vRealize Operations Manager master node.
- The vRealize Operations Manager master replica node.
- All vRealize Operations Manager data nodes.
- All vRealize Operations Manager remote collector nodes.

- 7 Power on the vRealize Log Insight virtual appliances in the management domain.

Important You must wait until each VM or virtual appliance is powered on and all of its services running before powering on the next VM.

Power on the VMs in the following order:

- The vRealize Log Insight master node.

- All vRealize Log Insight worker nodes.
- 8 Power on the vRealize Suite Lifecycle Manager appliance in the management domain.
 - 9 Power on the VMs in the first VI workload domain.

Important Each workload domain includes a three-node NSX Controller cluster. Power on these VMs first.

- a On the SDDC Manager Dashboard, navigate to the management domain.
 - b In the domain details panel, click the launch for the vCenter Server instance.
A new browser tab opens and displays the vSphere Web Client.
 - c In the vSphere Web Client, power on the VMs in the following order:
 - The three-node NSX Controller cluster.
 - The workload domain VMs.
 - d Repeat this procedure on each VI workload domain.
- 10 Power on the VMs in the Horizon domains (if any).
 - a On the SDDC Manager Dashboard, navigate to the first Horizon domain.
 - b In the domain details panel, launch the vCenter Server instance.
A new browser tab opens and displays the vSphere Web Client.
 - c In the vSphere Web Client, power on the Horizon domain VMs.
 - d Repeat this procedure on each Horizon domain.
 - 11 Using SSH, log in to the SDDC Manager VM as **vcf**.
 - 12 Run the following command:

```
sudo /opt/vmware/sddc-support/sos --health-check
```
 - 13 When prompted, enter the **vcf** user password.
Verify that everything works correctly.

Replace Host Components

21

The replacement procedure depends on the component being replaced and the condition of the component.

- [Replacing Components of a Host Running in Degraded Mode](#)

The procedures for replacing components of hosts in degraded depend on whether the host is part of a workload domain.

- [Replace a Dead Host](#)

If you need to replace a dead host, you must remove the host from the physical rack. You can then add a new host or replace the failed component on the host and add it back.

- [Replace Boot Disk on a Host](#)

This section describes the replacement procedure for a failed boot disk on a host.

Replacing Components of a Host Running in Degraded Mode

The procedures for replacing components of hosts in degraded depend on whether the host is part of a workload domain.

These procedures apply to the following components:

- CPU
- Memory
- BMC
- Power supply
- RAID 1 boot disk

Replace Components of a Workload Domain Host Running in Degraded Mode

This procedure shows you how to replace the component of a degraded host that is part of a workload domain.

Prerequisites

- Verify that the host is operational and is accessible by VMware Host Client.
- Verify that the Management, vSAN, and vMotion networks are available on the host. This can be viewed through the **Inventory > Hosts** page.
- Verify that the HDD and SSD disks on the host are in a good state.

Procedure

- 1 Log in to vSphere Web Client.
- 2 Right-click the affected host and click **Enter Maintenance Mode**.
- 3 If the host belongs to a domain, click **Full Data Migration**.
- 4 Right-click the affected host and select **Shutdown**.
- 5 Pull the host out of the physical rack.
Note the ports on the switches it was connected to.
- 6 Service the appropriate part following the OEM vendor documentation.
- 7 Put the host back in the physical rack and connect it back to the appropriate switches.
- 8 Power on the host.
- 9 In vSphere Web Client, right-click the host and click **Exit Maintenance Mode**.

Replace Components of an Unassigned Host Running in Degraded Mode

This procedure shows you how to replace the component of a degraded host that is not part of a workload domain.

Prerequisites

- Verify that the host is operational and is accessible by VMware Host Client.
- Verify that the HDD and SSD disks on the host are in a good state.

Procedure

- 1 Log in to vSphere Web Client.
- 2 Right-click the affected host and select **Shutdown**.
- 3 Pull the host out of the physical rack.
Note the ports on the switches it was connected to.
- 4 Service the appropriate part following the OEM vendor documentation.
- 5 Put the host back in the physical rack and connect it back to the appropriate switches.
- 6 Power on the host.

- 7 In the SDDC Manager Dashboard, verify that the host is available in the free pool.

Replace a Dead Host

If you need to replace a dead host, you must remove the host from the physical rack. You can then add a new host or replace the failed component on the host and add it back.

This procedure applies chiefly to the following components:

- Storage controllers
- Motherboards
- Boot disks

Prerequisites

If the host belongs to a workload domain, verify that there are at least 4 hosts in the management or workload domain to which the faulty host belongs. If there are less than 4 hosts, add a host to the domain from the capacity pool if possible.

Procedure

- 1 Decommission the host.
See [Decommission Hosts](#).
- 2 Power off the host and remove it from the physical rack.
- 3 Replace and reconfigure, as follows.
 - a Replace the failed component on the host.
 - b Perform a fresh reinstall of ESXi.
 - c Commission the host.
See [Commission Hosts](#).

Replace Boot Disk on a Host

This section describes the replacement procedure for a failed boot disk on a host.

Prerequisites

Verify that there are at least 4 hosts in the management or workload domain to which the faulty host belongs. If there are less than 4 hosts, add a host to the domain from the capacity pool, if possible.

Procedure

- 1 If there are dual boot disks in the host setup as RAID 1 and only one of them fails:
 - See [Replacing Components of a Host Running in Degraded Mode](#) to replace the failed disk.

The RAID 1 feature will rebuild the disks as needed. For more details, refer to the OEM vendor documentation.

- 2 If there is a single boot disk in the host and it fails, see [Replace a Dead Host](#).

User and Group Management

22

You can allow the users and groups in your Microsoft Active Directory (AD) domain to use their credentials to log in to the SDDC Manager Dashboard as well as the vCenter Server instances that are deployed in your Cloud Foundation system.

You had provided a password for the superuser account (user name vcf) in the deployment parameter sheet before bring-up. After Cloud Foundation is deployed, you can log in with the superuser credentials and then add vCenter Server or AD users or groups to Cloud Foundation. Authentication to the SDDC Manager Dashboard uses the VMware vCenter® Single Sign-On authentication service that is installed with the Platform Services Controller feature during the bring-up process for your Cloud Foundation system.

This chapter includes the following topics:

- [Assign Cloud Foundation Role to AD Users or Groups](#)
- [View Role Details](#)
- [Remove Cloud Foundation Role for a User or Group](#)

Assign Cloud Foundation Role to AD Users or Groups

You can assign the Cloud Admin role to AD users or groups so that they can log in to SDDC Manager with their AD credentials.

Procedure

- 1 Log in to the SDDC Manager Dashboard with your superuser credentials.
- 2 Click **Administration > Users**.
- 3 Click **+ User or Group**.
- 4 Select one or more user or group by clicking the check box next to the user or group.
You can either search for a user or group by name, or filter by user type or domain.
- 5 Scroll down to the bottom of the page and click **Add**.

The Cloud Admin role is assigned to the selected user or group.

View Role Details

The Cloud Admin role has read, write, and delete privileges.

Procedure

- 1 On the SDDC Manager, click **Administration > Users**.
- 2 In the Role column, click Cloud Admin.

The Role Details page displays privilege for the Cloud Admin role.

Remove Cloud Foundation Role for a User or Group

You can remove the Cloud Admin role from an AD user or group. The removed user or group will not be able to log in to the SDDC Manager Dashboard.

Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Users**.
- 2 Hover your mouse in the user or group row that you want to remove.
Three dots appear to the left of the user/group name column.
- 3 Click the dots and click **Remove User**.

The Cloud Admin role is removed for the specified user.

For security reasons, you can change passwords for the accounts that are used by your Cloud Foundation system. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

You specified passwords for your Cloud Foundation system as part of the bring-up procedure. You can rotate and update some of these passwords using the password management functionality in the SDDC Manager Dashboard or by using cURL API requests. For example:

- Accounts used for service consoles, such as the ESXi root account.
- The single sign-on administrator account.
- The default administrative user account used by virtual appliances.

To provide optimal security and proactively prevent any passwords from expiring, you should rotate passwords every 80 days.

Some tasks require dual authentication, that is, they required a privileged user name and password. You must configure and update the privileged user and password using the vSphere Client. See [Configure Dual Authentication](#).

Note Passwords for vRealize, Horizon 7, and PKS components cannot be managed through the SDDC Manager Dashboard.

This chapter includes the following topics:

- [Configure Dual Authentication](#)
- [Rotate Passwords for Managed Entities](#)
- [Manually Update Passwords](#)
- [Look Up Account Credentials](#)
- [Password Management cURL API Reference](#)
- [Updating SDDC Manager Passwords](#)

Configure Dual Authentication

You must configure dual authentication in order to perform certain tasks, such as updating or rotating passwords and configuring NSX Manager backups.

You will use the vSphere Client to create a new SSO group (`Sddc_Secured_Access`), add a user to the group, and assign a password to that user. The user is called the privileged user and will be required, along with its password, to perform certain tasks from the SDDC Manager UI or the VMware Cloud Foundation API.

You can create a new SSO user as the privileged user, or use an existing SSO user. If you plan to invoke operations requiring the privileged user as part of an automation solution, you should create a separate SSO user for this purpose. The SSO users used by automation should also be assigned the `No Access` role.

Note The `administrator@vsphere.local` user cannot be the privileged user.

Prerequisites

To perform this operation, you need to log in to the management vCenter Server as the `administrator@vsphere.local` user or another user who has the administrator role.

Procedure

- 1 Log into management vCenter Server using the vSphere Client.
- 2 Navigate to **Administration > Single Sign On > Users and Groups**.
- 3 Click the **Users** tab and select the domain from the drop-down list.
- 4 To create a new user in the selected domain, click **Add User**, enter the required information, and click **Add**.
- 5 Click the **Groups** tab and click **Add Group**.
- 6 Create a group named `Sddc_Secured_Access`, add the new or existing user to the group, and click **Add**.

Rotate Passwords for Managed Entities

As a security measure, you can rotate passwords for the logical and physical entities on all racks in your system. The process of password rotation generates randomized passwords for the selected accounts.

You can rotate passwords for the following entities.

- ESXi
- vCenter Server
 - By default, the vCenter Server root password expires after 90 days.
- PSC
 - By default, the PSC password expires after 90 days.
- NSX Manager
- NSX Controllers (NSX for vSphere and NSX-T)

- NSX Edge
- NSX-T Manager
- vRealize Log Insight
- vRealize Operations
- vRealize Automation
- vRealize Suite Lifecycle Manager
- SDDC Manager **backup** user

To update the SDDC Manager root, super user, and API passwords, see [Updating SDDC Manager Passwords](#).

Prerequisites

- Verify that there are no currently failed workflows in your Cloud Foundation system. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the brief time period that the password rotation process is running. It is recommended that you schedule password rotation for a time when you expect to have no running workflows.
- Configure the privileged user. For more information, see [Configure Dual Authentication](#).

Procedure

- 1 From the navigation pane, choose **Administration > Security > Password Management > Locally Managed..**

The Password Management page displays a table with detailed information about all domains, including their component, credential type, FQDN, IP address, and user name. This table is dynamic. Each column can be sorted.

You can click the filter icon next to the table header and filter the results by a string value. For example, click this icon next to **User Name** and enter **admin** to display only domains with that user name value.

- 2 Select the component type for which you want to rotate passwords from the **Component** drop-down menu. For example, **ESXI**.
- 3 Select one or more components and click **Rotate**.
- 4 Enter the privileged username and the privileged password and click **Rotate**.

A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status for the password rotation operation. Click on the task name to view sub-tasks. As each of these tasks are run, the status is updated. If the task fails, you can click **Retry**.

Results

Password rotation is complete when all sub-tasks are completed successfully.

Manually Update Passwords

You can manually change the password for a selected domain account. Unlike password rotation, which generates a randomized password, you provide the new password.

You can modify only one password at a time.

Note You cannot use these controls to update the NSX-T password. You can only update the NSX-T password from the NSX-T Manager product interface.

Prerequisites

- Verify that there are no currently failed workflows in your Cloud Foundation system. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the manual password update.
- Configure the privileged user. For more information, see [Configure Dual Authentication](#).

Procedure

- 1 From the navigation pane, choose **Administration > Security > Password Management**.

The Password Management page displays a table with detailed information about all domains, including their component, credential type, FQDN, IP address, and user name. This table is dynamic. Each column can be sorted.

You can click the filter icon next to the table header and filter the results by a string value. For example, click this icon next to **User Name** and enter **admin** to display only domains with that user name value.

- 2 Select the domain entity whose password you want to update and click **Update** at the top of the page.

Note If you select more than one domain, the **Update** button is disabled.

The Update Password dialog box appears. This dialog box also displays the entity name, credential type, user name, privileged user name, privileged password in case you need to confirm you have selected the correct domain. Enter the values for all these fields.

- 3 Enter the privileged username and privileged password.
- 4 Enter and confirm the new password.

If the passwords do not match, the dialog displays a red alert.

5 Click **Update**.

A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status of the password update operation. Click on the task name to view sub-tasks.

If the Tasks panel shows the task as having failed, click **Retry**.

Results

Password update is complete when all sub-tasks are completed successfully.

Look Up Account Credentials

To look up the account credentials for the built-in accounts that are managed and rotated by SDDC Manager, you log in to the SDDC Manager VM using the root account credentials.

Prerequisites

You must have the root account credentials to log in to the SDDC Manager VM.

Configure the privileged user. For more information, see [Configure Dual Authentication](#).

Procedure

- 1 SSH in to the SDDC Manager VM using the **vcf** user account.
- 2 (Optional) Change to the `/usr/bin` directory.

Note Although the password management CLI commands are located in `/usr/bin`, you can run them from any directory.

- 3 Enter `su` to switch to the root user.
- 4 Obtain the account credentials list by typing the command:

```
lookup_passwords
```

You will be required to enter the privileged user name and the privileged password.

To display the output in JSON format, use the following example command:

```
curl "https://localhost/security/password/vault" -k -u "<administrative user name>:<password>" -H "Accept: application/json" -H "privileged-username: vcf-secure-user@vsphere.local" -H "privileged-password: AfGh!8f9"
```

Enter the required credentials.

- 5 (Optional) Save the command output to a secure location with encryption so that you can access it later and use it to log in to the components as needed.

Password Management cURL API Reference

You can perform basic password management operations using cURL API requests. SSH in to the SDDC Manager VM and log in as the root user to use the cURL API.

cURL Password Operation API Requests

Some of the above operations can be run using cURL API requests.

Look up passwords - JSON format

Retrieves and lists in JSON format the account credentials for the built-in accounts that are managed and rotated by SDDC Manager.

```
# curl 'http://localhost/security/password/vault' \
-i -H 'Accept: application/json'
```

Look up passwords - plain text format

Retrieves and lists in plain text format the account credentials for the built-in accounts that are managed and rotated by SDDC Manager.

```
# curl 'http://localhost/security/password/vault' \
-i -H 'Accept: text/plain'
```

Update password

Updates the password of the specified domain component.

```
# curl 'http://localhost/security/password/vault' -i -X POST \
-H 'Content-Type: application/json' \
-H 'Accept: application/json' \
-d '{
  "entities": [{
    "credentialType" : "<credential type such SSH or API>",
    "entityIpAddress" : "<IP address>",
    "entityType" : "<component, such as ESXI>",
    "entityId" : "<node ID value>",
    "password" : "<password>",
    "domainName" : "<domain name>",
    "entityName" : "<FQDN>",
    "username" : "root"
  }],
  "type": "UPDATE"
}'
```

Rotate password

Rotates the password of the specified domain component.

```
# curl 'http://localhost/security/password/vault' -i -X POST \
-H 'Content-Type: application/json' \
-H 'Accept: application/json' \
-d '{
  "entities": [{
    "credentialType" : "<credential type such SSH or API>",
    "entityIpAddress" : "<IP address>",
    "entityType" : "<component, such as ESXI>",
    "entityId" : "<node ID value>",
    "password" : "<password>",
    "domainName" : "<domain name>",
    "entityName" : "<FQDN>",
    "username" : "root"
  }],
  "type": "ROTATE"
}'
```

Password operation history

Returns in JSON format the password history recorded in the password management database.

```
# curl 'https://localhost/security/password/vault/transactions' \
-i -H 'Accept: application/json' \
-k -u "<administrative user name>:<password>"
```

Password operation status

Returns in JSON format the latest (or current) workflow, which is an asynchronous job running in SDDC Manager. It polls the status of the workflow and reports percentage completed until the workflow finishes, at which time it reports its status.

```
# curl 'https://localhost/security/password/vault/transactions/2002' \
-i -H 'Accept: application/json' \
-k -u "<administrative user name>:<password>"
```

Retry failed password operation

Retries the specified failed operation and returns results in JSON format

```
# curl 'http://localhost/security/password/vault/transactions/2002' \
-i -X PATCH \
-H 'Content-Type: application/json' \
-H 'Accept: application/json' \
-d '{
  "entities": [{
    "credentialType" : "<credential type such SSH or API>",
    "entityIpAddress" : "<IP address>",
    "entityType" : "<component, such as ESXI>",
    "entityId" : "<node ID value>",
    "password" : "<password>",
  }],
  "type": "RETRY"
}'
```

```

    "domainName" : "<domain name>",
    "entityName" : "<FQDN>",
    "username" : "root"
  }],
  "type": "<specify ROTATE or UPDATE>"
}'

```

Cancel password operation

Cancels failed password operations and returns results in JSON format

```

# curl 'https://localhost/security/password/vault/transactions/2002' \
-i -X DELETE -H 'Accept: application/json' \
-k -u "<administrative user name>:<password>"

```

Updating SDDC Manager Passwords

You cannot update SDDC Manager passwords through the SDDC Manager Dashboard or by using cURL API requests. Instead, you will need to SSH into the SDDC Manager VM and make the changes there.

The process for updating SDDC Manager passwords varies, depending on which account you are updating.

- [Update SDDC Manager Root and Super User Passwords](#)

For security reasons, you can change passwords for the SDDC Manager root (**root**) and super user (**vcf**) accounts. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

- [Update SDDC Manager REST API Account Password](#)

To use the VMware Cloud Foundation API, an API client logs in using the SDDC Manager **admin** account. For security reasons, you should periodically update the password for this account.

- [Update Expired SDDC Manager root Password](#)

This section describes the procedure for updating an expired password for the SDDC Manager root (**root**) user.

Update SDDC Manager Root and Super User Passwords

For security reasons, you can change passwords for the SDDC Manager root (**root**) and super user (**vcf**) accounts. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

The SDDC Manager **root** password expires after 365 days.

Procedure

- 1 SSH in to the SDDC Manager VM using the **vcf** user account.

- 2 Enter **su** to switch to the root user.
- 3 Enter one of the following commands:

Option	Description
<code>passwd vcf</code>	To change the super user password.
<code>passwd root</code>	To change the root password.

- 4 Enter and retype the new password. For example:

```
root@sddc-manager [ /home/vcf ]# passwd vcf
New password:
Retype new password:
passwd: password updated successfully
```

Results

The password is updated.

Update SDDC Manager REST API Account Password

To use the VMware Cloud Foundation API, an API client logs in using the SDDC Manager **admin** account. For security reasons, you should periodically update the password for this account.

If you write a script that invokes the APIs, the script should either prompt the user for the password for the **admin** account or should accept the password as a command line option. As a best practice, you should not encode the password for the account in the script code itself.

Password requirements:

- Length 8-12 characters
- Must include: mix of upper-case and lower-case letters a number a special character such as @ ! # \$ % ^ or ?
- Cannot include: * { } [] () / \ ' " ` ~ , ; : . < >

Procedure

- 1 SSH in to the SDDC Manager VM using the **vcf** user account.
- 2 Enter **su** to switch to the root user.
- 3 Enter the following command:

```
/opt/vmware/vcf/commonsvcs/scripts/auth/set-basicauth-password.sh admin <password>
```

For *<password>*, enter the new password for the **admin** account.

Update Expired SDDC Manager root Password

This section describes the procedure for updating an expired password for the SDDC Manager root (**root**) user.

Prerequisites

Procedure

- 1 Log into the vSphere Client and select the SDDC Manager VM from Mgmt-ResourcePool.
- 2 From the panel on the right side of the window, select **Summary** and click on **Launch Web Console**.
- 3 In the popup window, select web console and click **OK**. This will open the console in a new browser tab.

The console opens in a new browser window.

- 4 Click **Login**.
- 5 Type **root** as the user name and enter the current password for the root user.
- 6 When prompted for current password, enter the current password.
- 7 When prompted for a new password, enter a different password than the previous one and click **OK**.

Backing Up and Restoring SDDC Manager

24

Back up the SDDC Manager VM regularly to avoid downtime and data loss in case of a system failure. If the SDDC Manager VM does fail, you can restore VM to the last backup.

Follow the best practises below:

- Schedule backups when SDDC Manager is not running any workflows.
- Take periodic backups on a daily to weekly frequency.
- If a workflow does not complete successfully and the Cloud Foundation environment is in this state when the scheduled backup is taken, resolve the failure as soon possible and take an unscheduled backup. Restoring your environment from a backup that includes unresolved failures is more difficult than restoring from a clean backup.

A workflow is resolved when the environment is not in an intermediate state. For some SDDC Manager workflows, the workflow can only be resolved by fixing the failure conditions and retrying the operation. Other workflows can also be resolved by invoking the corresponding delete operation. For example, if adding a host to a workload domain fails, either fix the condition that caused the workflow to fail, or run the workflow that removes the host from the cluster. Contact VMware support if you are unable to resolve a workflow.

You can back up and restore SDDC Manager with an image-based or a file-based solution (starting from Cloud Foundation 3.7.2). It is recommended that you use the image-based solution. The file-based approach is more limited, has a higher network-ingress cost since patch and install bundles have to be re-downloaded, and requires scripting and working with APIs.

For a file-based backup of the SDDC Manager VM, you must manually configure the NSX Managers in your VMware Cloud Foundation environment to back up their state to an FTP site. You must repeat this step every time you create a workload domain.

This chapter includes the following topics:

- [Image-Based Backup and Restore](#)
- [File-Based Backup and Restore](#)

Image-Based Backup and Restore

For an image-based backup of the SDDC Manager, use a solution compatible with the VMware vSphere Storage APIs - Data Protection (formerly known as VMware vStorage APIs for Data Protection or VADP).

vSphere Storage APIs - Data Protection compatible backup software connects to the vCenter servers in the management domain to perform the backup to a remote site. In the event of failure, the backup software connects to the vCenter servers in the management domain to restore the VMs. If the management domain is lost, the vCenter servers are no longer available and must be restored first. Choosing a backup software that supports Direct Restore to an ESXi host allows restoring the vCenter servers.

For an SDDC Manager backup, connect your backup with the management domain vCenter Server. Configure the product to take non-quieted backups of SDDC Manager. To reduce the backup time and storage cost, use incremental backups in addition to full ones.

File-Based Backup and Restore

With Cloud Foundation 3.7.2 and later, you can use APIs for a file-based backup and restore solution for SDDC Manager. The APIs are building blocks and do not implement a complete solution.

Before getting into the SDDC Manager details, it is useful to review the basic principles in a file-based solution. In such a solution, the state of a product is periodically exported to a file that is stored in a fault domain different than the one where the product is running. If the product needs to be restored, the OVA is redeployed and a selected backup file is used to restore the state. Finally, the post-restore steps are done.

In case you have to restore the SDDC Manager VM, you select the backup file to restore and download the appropriate OVA file. You can deploy this OVA either through vCenter Server or the OVF tool. You then load the state on the newly deployed SDDC Manager VM.

Note the following limitations for a file-based backup of the SDDC Manager VM:

- This solution requires that you register an external SFTP server. See . If you do not use an external SFTP server, the NSX Managers continue to write the backups to the built-in SFTP server, and so NSX managers are not protected. They are not protected because the SDDC Manager file-based backup mechanism does not back up the NSX Manager backup files.
- This solution cannot be used for composable servers.
- This solution cannot be used when you have stretched clusters in your environment.

Configure an External SFTP Server for NSX Manager Backups

VMware Cloud Foundation allows you to register an external SFTP server with SDDC Manager for backing up NSX Managers.

Until you register an external SFTP server, the NSX backups are taken on the SFTP server that is built into SDDC Manager. It is recommended that you register an external SFTP server soon after you upgrade or deploy VMware Cloud Foundation. Using an external SFTP server provides better protection against failures because it decouples the NSX backups from the SDDC Manager backups. The built-in SFTP server provides temporary protection against failures and should be used while you are setting up an external SFTP server.

It is important to deploy a reliable SFTP server and ensure it is accessible from the VMware Cloud Foundation instance. If the SFTP server is not available when an NSX Manager attempts to back up its state, the backup will not be taken, and any recent changes are not backed up until the retries succeed. To ensure that this situation does not occur, it is recommended that you periodically check that NSX Manager backups are successfully taken, and monitor that the backups for other products are also being successfully taken. If the SFTP server is not available at the time of deploying a workload domain or upgrading NSX, these operations fail.

When you register an SFTP server with SDDC Manager, it saves the SFTP server details, and then configures all existing NSX Managers to use the SFTP server. Finally, when any subsequent NSX Managers are deployed, SDDC Manager configures them to use this SFTP server as well. When you register the SFTP server, you must also specify a phrase to use to encrypt the NSX Manager backups. Note that this same phrase is also used to encrypt SDDC Manager file-based backups. For more information, see [Chapter 24 Backing Up and Restoring SDDC Manager](#). You can use the same UI and API to edit the settings of an already configured SFTP server and encryption phrase.

To configure an external SFTP server for the NSX Manager backup, perform the following steps:

Note The backup server is available only for NSX Manager whereas the passphrase is available for both SDDC Manager and NSX Manager.

Prerequisites

- The external SFTP server must support ECDSA SSH public key.
- You must configure a privileged user and password before you can configure an external SFTP server. See [Configure Dual Authentication](#).

Procedure

- 1 In the SDDC Manager dashboard, select **Administration > Backup Configuration**.
- 2 Click **+Register External**.
- 3 Enter the IP address of the backup server. Ensure that the server is available for the successful configuration.
- 4 Enter the port number at which the SFTP service is running.
- 5 Enter the credentials of the server.

- 6 Enter the backup directory path of the server. Ensure that the user you specify in step 5 can access the directory path since the backups are saved to this location. It is recommended to provide different directory paths for the different VMware Cloud Foundation instances in case you are using the same SFTP server across all.
- 7 Confirm the fingerprint that is auto populated for the given IP address and the port.
- 8 Enter the passphrase which is used for both NSX Managers and SDDC Manager backups.
- 9 In the Authentication Credentials section, enter the privileged user credentials.

The privileged user has access to privileged data. You created this user when you configured dual authentication.
- 10 Click **Save**.
- 11 Click **Confirm**.
- 12 If you have to edit the backup configuration information, perform the following steps:
 - a On the SDDC Manager dashboard, select **Administration > Backup Configuration**.
 - b Click **Edit**.
 - c Edit the text boxes as per your requirement. If there is any change in the IP address or the backup directory path, if you save the configuration, the existing backups are not copied to the new location. Copy them manually.
 - d Enter the backup server password and the passphrase. If there is any change in the passphrase to the existing, you need to use the old passphrase while restoring previously taken backups.
 - e Enter the privileged user credentials.
 - f Click **Save**.
 - g Click **Confirm**.

Backup SDDC Manager

See Section 2.14 in the *VMware Cloud Foundation API Reference Guide* for the manual procedure to take the backup by using SDDC Manager APIs. It is recommended that you write a script to automate this process.

Restore SDDC Manager

See Section 2.14 of the 3.9 API guide for the manual procedure to restore SDDC Manager from file-based backups using the SDDC Manager APIs.

Cloud Foundation Glossary

25

Term	Description
availability zone	Collection of infrastructure components. Each availability zone is isolated from other availability zones to prevent the propagation of failure or outage across the data center.
bring-up	Initial configuration of a newly deployed Cloud Foundation system. During the bring-up process, the management domain is created and the Cloud Foundation software stack is deployed on the management domain.
commission host	Adding a host to Cloud Foundation inventory. The host remains in the free pool until it is assigned to a workload domain.
composability	Ability to dynamically configure servers to meet the needs of your workloads without physically moving any hardware components. You bind disaggregated hardware components (compute, network, storage, and offload components) together to create a logical system based on the needs of your applications.
dirty host	A host that has been removed from a cluster in a workload domain. A dirty host cannot be assigned to another workload domain until it is cleaned up.
decommission host	Remove an unassigned host from the Cloud Foundation inventory. SDDC Manager does not manage decommissioned hosts.
free pool	Hosts in the Cloud Foundation inventory that are not assigned to a workload domain
host	An imaged server.
inventory	Logical and physical entities managed by Cloud Foundation.
Lifecycle Manager (LCM)	Automates patching and upgrading of the software stack.
management domain	Cluster of physical hosts that contains the management component VMs
network pool	Automatically assigns static IP addresses to vSAN and vMotion vmkernel ports so that you don't need to enter IP addresses manually when creating a VI workload domain or adding a host or cluster to a workload domain.
patch update bundle	Contains bits to update the appropriate Cloud Foundation software components in your management or VI workload domain.
region	A Cloud Foundation instance.
SDDC Manager	Software component that provisions, manages, and monitors the logical and physical resources of a Cloud Foundation system.
SDDC Manager VM	Virtual machine (VM) that contains the SDDC Manager services and a shell from which command line tools can be run. This VM exposes the SDDC Manager UI.
server	Bare metal server in a physical rack. After imaging, it is referred to as a host.

Term	Description
unassigned host	Host in the free pool that does not belong to a workload domain.
workload domain	A policy based resource container with specific availability and performance attributes that combines vSphere, storage (vSAN, NFS, or VMFS on FC) and networking (NSX for vSphere or NSX-T) into a single consumable entity. A workload domain can be created, expanded, and deleted as part of the SDDC lifecycle operations. It can contain cluster(s) of physical hosts with a corresponding vCenter to manage them. The vCenter for a workload domain physically lives in the management domain.